

IBM Tivoli Monitoring
Version 6.3 Fix Pack 1

Installation and Setup Guide



IBM Tivoli Monitoring
Version 6.3 Fix Pack 1

Installation and Setup Guide



Note

Before using this information and the product it supports, read the information in “Notices” on page 963.

This edition applies to version 6, release 3, fix pack 1 of IBM Tivoli Monitoring (product number 5724-C04) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2005, 2013.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	xv
--------------------------	-----------

Tables	xix
-------------------------	------------

Chapter 1. Introduction 1

New in this release	1
New in Version 6.3 Fix Pack 1	1
New in Version 6.3	1
Overview of IBM Tivoli Monitoring	5
Components of the monitoring architecture	6
Tivoli Enterprise Monitoring Server	10
Tivoli Enterprise Portal	11
Tivoli Enterprise Monitoring Agents	11
Tivoli Data Warehouse.	13
Event synchronization component	13
Tivoli Enterprise Portal Server extended services	14
tacmd Command Line Interface	14
IBM Dashboard Application Services Hub and	
dashboard applications	15
Tivoli Authorization Policy Server	15
tivcmd Command Line Interface for	
Authorization Policy	16
Tivoli Enterprise Monitoring Automation Server	17
Tivoli Performance Analyzer.	17

Chapter 2. Pre-deployment phase 19

Planning checklist	19
Understanding Tivoli Monitoring and your network	21
Determine if you require a firewall gateway	22
Determine where to place your Tivoli Monitoring	
components	23
Additional ports used in the Tivoli Monitoring	
environment	34
Sizing your Tivoli Monitoring hardware	34
Locating and sizing the hub Tivoli Enterprise	
Monitoring Server	35
Locating and sizing the remote Tivoli Enterprise	
Monitoring Server	35
Locating and sizing the remote deployment	
depot	36
Locating and sizing the Tivoli Enterprise Portal	
Server	37
Locating and sizing the Warehouse Proxy Agent	37
Locating and sizing the Summarization and	
Pruning Agent	39
Locating and sizing the portal client	40
Locating and sizing the Tivoli Performance	
Analyzer	40
Using IBM Installation Manager	42
Software Product Compatibility Reports (SPCR)	43
Configuring for high availability and disaster	
recovery	43
Configuring the hub monitoring server for high	
availability and disaster recovery	43

Configuring for portal server high availability	
and disaster recovery	44
Configuring for IBM Dashboard Application	
Services Hub high availability and disaster	
recovery	45
Configuring for agent and remote monitoring	
server high availability and disaster recovery	45
Configuring for warehouse high availability and	
disaster recovery	46
Configuring for Warehouse Proxy Agent high	
availability and disaster recovery	46
Configuring for Summarization and Pruning	
Agent high availability and disaster recovery	46
Configuring for Tivoli Performance Analyzer	
high availability and disaster recovery	47
Agent deployments.	47
Self-describing monitoring agents	48
Prerequisite Scanner for IBM Tivoli Monitoring	
agents	48
Background information about agent autonomy	55
Agentless monitoring versus monitoring agents	56
Mainframe users.	62
Multi-hub environments	63
Accelerating your custom monitoring.	65
Planning and project management.	66
Estimating deployment tasks	66
Install server components on Windows and	
UNIX	67
Install server components on z/OS	67
Install data warehousing components.	67
Install and configure event integration	
components	68
Install and configure monitoring agents	68
Setting up LDAP and single sign-on for your	
dashboard environment	69
Creating authorization policies	70
Setting up situation-based monitoring	70
Creating workflow policies	70
Creating workspaces	70
Creating and deploying Agent Builder Agents.	70
Transferring skills	71
Scheduling the initial deployment	71
Scheduling for fix packs	71
Staffing.	72

Chapter 3. Hardware and software requirements 75

Supported operating systems	75
Supported databases for Tivoli Enterprise Portal	
Server	82
Supported databases for Tivoli Data Warehouse	83
Required hardware for distributed systems	84
Processor requirements	84
Memory and disk requirements.	85
Additional requirements	87

Required hardware for System z	88
Required software	89
Supported browser versions	90
Required software for event integration with Netcool/OMNIbus	91
Required software and memory requirements for a dashboard environment	91
Required software and memory requirements for a reporting environment that uses Tivoli Common Reporting	93
Required software and memory requirements for using the OSLC Performance Monitoring service provider	94
Required software and memory requirements for non-linear trending in Tivoli Performance Analyzer	95

Chapter 4. Deployment phase 97

Pre-installation checklist	97
Installing the infrastructure components	97
Configuration checklist	98
Customizing your environment	99
Installing your first 50 agents	100
Postinstallation checklist	101
Configuring your warehouse agents	101
Installing additional agents	102

Chapter 5. Tivoli Monitoring Startup Center 103

Startup Center platform support	103
Startup Center prerequisite requirements	104
Deployment procedure	105
Local installation	105
Distributed installation	108
Default values taken by Startup Center components	114
Remote Execution and Access	115
Windows Targets	115
Disabling User Account Control to facilitate RXA	116

Chapter 6. Post-deployment phase 119

Applying maintenance	119
Planning an upgrade	119
Upgrade steps	119
Post-upgrade health check	120
Maintaining an efficient monitoring environment	121
Daily health checks	122
Weekly health checks	123
Monthly health checks	123
Quarterly health checks	123

Chapter 7. Preparing for installation 125

Specific information to have ready	126
Information to gather for event forwarding	126
Naming your monitoring server	127
Choose between IPv6 and IPv4	127
Required order of installation or upgrade of IBM Tivoli Monitoring component products	128
Windows installation considerations	129
Windows 32 bit On Windows 64 bit (WOW64)	129

32 bit versus 64 bit	129
Installation using a Citrix client	130
Memory leak issue on Windows Server 2008	130
Linux or UNIX installation considerations	130
Changes in the behavior of the autostart scripts	130
Create an IBM Tivoli account for installing and maintaining the installation directory	134
Host name for TCP/IP network services	134
Use of fully qualified path names	135
Multiple network interface cards	135
Installing into an NFS environment	135
Installing into Solaris zones	135
Architecture and product codes	136
File descriptor (maxfiles) limit on UNIX and Linux systems	137
Installing into an existing installation	137
Security options	137
Communication between components	138
Authorization and authentication	139
Single sign-on capability	141
SOAP server security	142
Global Security Toolkit	142

Chapter 8. Upgrading from a previous installation 145

Planning your upgrade	145
Platforms no longer supported for IBM Tivoli Monitoring V6.3	145
Prerequisites for IBM Tivoli Monitoring V6.3	146
Components to upgrade	149
Required order of upgrade	150
Migrated information when upgrading from a previous version	152
Backing up IBM Tivoli Monitoring	152
Backing up your portal server and Tivoli Data Warehouse databases	155
Upgrading the warehouse from IBM Tivoli Monitoring V6.1.x levels	156
IBM Tivoli Monitoring coexistence and interoperability	158
Tivoli Enterprise Monitoring Server	158
Tivoli Enterprise Portal Server	159
Tivoli Data Warehouse	159
Agents	160
Tivoli Event Synchronization component	160
IBM Tivoli Monitoring V5.x interoperability	160
Upgrading from IBM Tivoli Monitoring V6.1 or V6.2	161
Overview of the upgrade process	161
Linux and UNIX: Upgrading a portal server running as a non-root process	163
Required Java Runtime Environment	165
Scenario: a rolling product upgrade	165
Configuration	165
Upgrading the Tivoli Monitoring environment	166
Expected results	167
Special instructions for reseeding a Hot Standby monitoring server	168

Chapter 9. Installing IBM Tivoli

Monitoring on one computer 171

Prerequisites for the single-computer installation 171

Installation procedure 172

Chapter 10. Installing IBM Tivoli

Monitoring 193

High-level installation steps 194

Installing and configuring the hub Tivoli Enterprise

Monitoring Server 196

Windows: Installing the hub monitoring server 196

Linux or UNIX: Installing the hub monitoring

server 200

Configuring the hub monitoring server. 202

Enabling self-describing agent capability at the

hub monitoring server 204

Dynamically controlling the hub monitoring

server self-describing agent capability 205

Configuring self-describing agent seeding 206

Default installation seeding values 207

Complete the following steps to enable

self-describing agent capability 208

Adding application support to the hub

monitoring server 208

Installing and configuring the remote monitoring

servers 211

Windows: Installing a remote monitoring server 211

Linux or UNIX: Installing a remote monitoring

server 215

Installing and configuring Tivoli Enterprise

Monitoring Automation Server 219

Windows: Installing a Tivoli Enterprise

Monitoring Automation Server 220

Linux or UNIX: Installing a Tivoli Enterprise

Monitoring Automation Server 222

Configuring the Tivoli Enterprise Monitoring

Automation Server 224

Installing the Tivoli Enterprise Portal Server 226

Windows: Installing the portal server 227

Linux or AIX: Installing the portal server 241

Installing monitoring agents 261

Windows: Installing a monitoring agent 262

Linux or UNIX: Installing a monitoring agent 267

Populating the data warehouse's

ManagedSystem table 271

Installing the Tivoli Enterprise Portal desktop client 271

Windows: Installing the desktop client 272

Linux: Installing the desktop client 273

Installing and enabling application support 275

Selecting the correct support media 277

Configuring application support on the IBM

Tivoli Monitoring Base DVDs 278

Configuring application support for nonbase

monitoring agents 279

Configuring application support on nonlocal

monitoring servers 292

Using IBM Installation Manager to install or

update IBM Tivoli Monitoring components for a

dashboard environment 299

Using the Installation Launchpad to install or

update IBM Tivoli Monitoring components for a

dashboard environment 299

Using IBM Installation Manager to install or

update components from a local repository 300

Using IBM Installation Manager to install or

update components from a network repository 302

Saving package information for rollback 305

Installing and configuring the IBM Infrastructure

Management Dashboards for Servers 306

Installing the IBM Infrastructure Management

Dashboards for Servers using the graphical user

interface 308

Installing or updating the IBM Infrastructure

Management Dashboards for Servers using

console mode 311

Configuring the IBM Infrastructure Management

Dashboards for Servers package after

installation 314

Updating the IBM Infrastructure Management

Dashboards for Servers using the graphical user

interface 315

Rollback of the IBM Infrastructure Management

Dashboards for Servers 317

Installing and configuring the Tivoli Authorization

Policy Server and tivcmd Command-Line Interface

for Authorization Policy 318

Installing and configuring the Tivoli

Authorization Policy Server using the graphical

user interface 320

Installing and configuring the Tivoli

Authorization Policy Server using console mode 325

Configuring the Tivoli Authorization Policy

Server feature after installation 328

Verifying the Tivoli Authorization Policy Server

installation 329

Installing the tivcmd Command Line Interface

for Authorization Policy using the graphical

user interface 330

Installing the tivcmd Command Line Interface

for Authorization Policy using console mode 332

Updating the Tivoli Authorization Policy Server

using the graphical user interface. 334

Rollback of the Tivoli Authorization Policy

Server 336

Installing domain definitions for Tivoli

Performance Analyzer 336

Installing domain definitions in GUI mode 337

Installing the domain definitions in console

mode 338

Installing language packs 339

Silent installation of language packs 340

Uninstalling a language pack 342

Configuring clients, browsers, and JREs 342

Desktop clients 343

Browser clients 343

Java Web Start clients. 347

Installing and configuring IBM Java 7 348

Upgrade installation versus pristine installation 348

Windows: Installing IBM Java 7 348

Linux: Installing IBM Java 7 357

Specifying the browser used for online help . . .	363
Windows: Specifying the browser location . . .	363
UNIX and Linux: Specifying the browser location . . .	364
Web Start: Specifying the browser location . . .	365
Starting the Tivoli Enterprise Portal client . . .	366
Starting the desktop client . . .	366
Starting the browser client . . .	366
Using Web Start to download and run the desktop client . . .	367
Enabling tracing for the JRE . . .	367
Downloading and running the desktop client . . .	368
Manually creating a shortcut for the Web Start client . . .	369
Installing product maintenance . . .	370

Chapter 11. Deploying monitoring agents across your environment . . . 371

Populating your agent depot . . .	371
Populating the agent depot from the installation image . . .	372
Populating the agent depot with the tacmd addBundles command . . .	373
Managing your agent depot . . .	374
Sharing an agent depot across your environment . . .	375
Deploying OS agents . . .	375
Requirements for the tacmd createNode command . . .	376
Using the tacmd createNode command . . .	377
Deploying non-OS agents . . .	377
Deploying through the portal . . .	378
Deploying through the command-line . . .	378
Deploying Netcool/OMNIbus System Service Monitor (SSM) agents . . .	379
Installing an SSM agent . . .	379
Uninstalling an SSM agent . . .	380
Installing an SSM patch . . .	380
Uninstalling an SSM patch . . .	380
Starting an SSM agent . . .	380
Stopping an SSM agent . . .	380
Restarting an SSM agent . . .	380
Configuring an SSM agent . . .	381
Bulk deployment of NetCool SSM agents . . .	381
Query deployment status of Netcool SSM agents . . .	381
Bulk agent deployment . . .	382
Deployment processing model . . .	382
Deploy status . . .	384
Organizing deployments using groups . . .	385
Best-practice deployment procedures . . .	389
Working with non-agent bundles . . .	392
Deploying a non-agent bundle . . .	392
Updating a non-agent bundle . . .	393
Removing a non-agent bundle . . .	393
Running deployment in a Hot Standby environment . . .	394
Self-describing agent installation . . .	394
Managing your self-describing agent installations . . .	394
Terminal self-describing agent installation errors . . .	396
Enabling or disabling self-describing agent capability . . .	396

Chapter 12. Monitoring your operating system via a System Monitor Agent. . 399

Installing the System Monitor Agent on Windows systems . . .	400
Configuring the System Monitor Agents on Windows . . .	402
Uninstalling the Windows System Monitor Agent . . .	403
Installing the System Monitor Agent on Linux or UNIX systems . . .	404
Contents of the silent response file . . .	405
Configuring the System Monitor Agents on Linux or UNIX . . .	406
Uninstalling the Linux or UNIX System Monitor Agent . . .	408
Defining common configuration parameters: accessing centralized configuration information . . .	409

Chapter 13. Configuring IBM Tivoli Monitoring components . . . 411

Starting Manage Tivoli Enterprise Monitoring Services . . .	411
Starting Manage Tivoli Enterprise Monitoring Services on Windows computers . . .	411
Starting Manage Tivoli Enterprise Monitoring Services on Linux or UNIX computers . . .	411
Changing the configuration of the Tivoli Enterprise Monitoring Server . . .	412
Configuring or changing the Tivoli Enterprise Monitoring Automation Server connection . . .	414
Windows: Configuring a Tivoli Enterprise Monitoring Automation Server . . .	414
Linux/UNIX: Configuring a Tivoli Enterprise Monitoring Automation Server . . .	415
Configuring or changing the monitoring server connection for agents . . .	417
Configuring the handling of pure events by the monitoring server . . .	419
Starting and stopping components . . .	420
Specifying network interfaces . . .	420
Controlling port number assignments . . .	421
Configuring port number assignments for the monitoring server . . .	421
Configuring port number assignments for the monitoring automation server . . .	421
Configuring port number assignments for the portal server . . .	421
Configuring port number assignments for monitoring agents . . .	423
Configuring the heartbeat interval . . .	426
Restarting the Tivoli Enterprise Portal Server after reconfiguration . . .	427
Verifying the dashboard data provider is enabled . . .	427
Switching to a different Tivoli Enterprise Portal Server database . . .	430
Silent configuration of the Performance Analyzer . . .	431
Configuring historical data collection for the Performance Analyzer warehouse agent . . .	432
Tivoli Monitoring protocol usage and protocol modifiers . . .	435

KDE_TRANSPORT Structure	435
KDE_TRANSPORT Transmission Control Protocol	435
Chapter 14. Additional Linux and UNIX configuration steps	439
Disabling fsync() calls	439
Configuring permissions for a monitoring server on a non-NIS Solaris computer	439
Linux requirements for the localhost host name	440
Setting ulimit values for the Warehouse Proxy Agent	440
Chapter 15. Configuring IBM Tivoli Monitoring components for IPv6 communication	441
Configuring AIX, Solaris, HP-UX, and Linux monitoring components	441
Configuring the hub monitoring server	441
Configuring the Tivoli Enterprise Monitoring Automation Server	443
Configuring the remote monitoring server	443
Configuring the monitoring agents	444
Warehouse Proxy Agent	445
Tivoli Enterprise Portal Server	446
Adding Application Support to a monitoring server	446
Command-line interface (CLI)	447
Configuring Windows monitoring components	447
Configuring the hub monitoring server	447
Configuring the Tivoli Enterprise Monitoring Automation Server	448
Configuring the remote monitoring server	448
Configuring the monitoring agents	448
Warehouse Proxy Agent	449
Tivoli Enterprise Portal Server	449
Adding Application Support to a monitoring server	449
Sending files to z/OS systems using FTP	450
Tivoli Enterprise Portal browser client	450
Tivoli Enterprise Portal desktop client	450
Command-line interface	450
Configuring z/OS monitoring components	451
Firewall considerations	451
Chapter 16. Additional Tivoli Enterprise Portal configurations	453
Connecting the Tivoli Enterprise Portal Server on Windows to a different monitoring server	453
Using SSL between the portal server and the client	454
Enabling and disabling SSL for the Tivoli Enterprise Portal Server	455
Disabling SSL	455
Reverting from the IBM HTTP Server to the internal web server	456
Portal server	456
Browser client	456
Desktop client	456
Web Start client	457

Configuring HTTP communication between the portal client and server	457
Browser client	457
Desktop client	458
Web Start client	458
Changing the ports used by IHS	459

Chapter 17. Configuring IBM Tivoli Monitoring Web Services (the SOAP Server)	461
Defining hubs	461
Windows: Defining hubs	461
UNIX and Linux: Defining hubs (GUI procedure)	462
UNIX and Linux: Defining hubs (CLI procedure)	463
Adding users	464
Windows: Adding users	465
UNIX or Linux: Adding users (GUI)	465
UNIX or Linux: Adding users (CLI)	466
Verifying the configuration	466

Chapter 18. Performance tuning	469
Disabling TCP-delayed acknowledgments on AIX systems	469
Tivoli Enterprise Monitoring Server	470
Tivoli Enterprise Monitoring Server tuning recommendations for large-scale environments	471
Tivoli Enterprise Monitoring agents	472
Agentless Monitoring	473
Adjust the thread pool size, refresh interval, and cache time-to-live parameters for your environment	473
To monitor many monitoring servers, spread the target monitoring servers across multiple process instances	473
Avoid collecting historical data for the KR2_WMI_Event_Log attribute group	474
Tivoli Enterprise Portal Server	474
Portal server parameter tuning	475
For environments with more than 10 events per minute, consider disabling the Event Assistant	478
Perform periodic table maintenance for Tivoli Enterprise Portal Server DB2 database tables	478
Tivoli Enterprise Portal client	478
Tuning the portal client JVM	478
Portal client parameter tuning	479
Avoid displaying a large number of events in the Situation Event Console	480
Tivoli Data Warehouse	480
For all but the simplest environments, database administrator skills are needed to manage the warehouse database	481
Historical data collection	481
Warehouse Proxy Agent	482
Summarization and Pruning Agent	486
Database tuning	489
Relational database design and performance tuning for DB2 Database servers	489
Optimizing queries	502

Processing queries.	502
Defining custom queries.	503
Optimizing situations.	504
Validating your installation.	506
CLI syntax	506
CLI example.	507

Chapter 19. Tivoli Data Warehouse solutions 509

New in Version 6.3	509
Planning considerations for the Tivoli Data Warehouse	510
Estimating the required size of your database	510
Understanding the disk requirements for your database	513
Increasing the size of your database (DB2 for Linux, UNIX, and Windows only)	514
Planning assumptions	515
Preliminary planning is complete.	515
Environment assumptions	516
The data warehouse is remote from the portal server	516
Agent warehouse database upgrade	516
Firewall considerations for the Tivoli Data Warehouse	517
Compressing historical data transmission	517
Next steps	519
Summary of supported operating systems.	519

Chapter 20. Schema Publication Tool 525

Generating SQL for data warehouse tables	525
Using the schema publication tool in updated mode	527
Using the schema publication tool in migrate mode	529
Using the Schema Publication Tool for database compression.	529

Chapter 21. Tivoli Data Warehouse range partitioning 531

Specifying range partitioned tables for the Warehouse Proxy Agent	532
Configuration parameters	533
Specifying range partitioned tables for the Summarization and Pruning Agent	533
Configuration parameters	534
Specifying range partitioned tables for the schema publication tool.	536

Chapter 22. Specifying where tables and indices are created 537

Environment variables	538
---------------------------------	-----

Chapter 23. Tivoli Data Warehouse solution using DB2 for Linux, UNIX, and Windows 539

Supported components	539
Prerequisite installation	541

Implementing a Tivoli Data Warehouse solution using DB2 for Linux, UNIX and Windows.	542
Solution steps	542

Step 1: Create the Tivoli Data Warehouse database 543

Creating the warehouse database on DB2 for Linux, UNIX and Windows.	544
Creating a warehouse user on Windows	544
Creating a warehouse user on Linux or UNIX	545
Limiting the authority of the warehouse user	545
Setting database and instance configuration values.	547
Activating the DB2 listeners on a UNIX DB2 server	547

Step 2: Install and configure communications for the Warehouse Proxy Agent 548

Cataloging a remote data warehouse	550
Configuring an ODBC data source for a DB2 data warehouse	551
Configuring a Warehouse Proxy Agent on Windows (ODBC connection)	552
Configuring a Warehouse Proxy Agent on Linux or UNIX (JDBC connection)	555
Starting the Warehouse Proxy Agent.	558

Step 3: Configure communications between the Tivoli Enterprise Portal Server and the data warehouse 559

Configuring a Windows portal server (ODBC connection)	559
Configuring a Linux or AIX portal server (DB2 for Linux, UNIX and Windows CLI connection)	561
Starting the portal server	562

Step 4: Install and configure communications for the Summarization and Pruning Agent 562

Step 5: Install and configure communications for the Tivoli Performance Analyzer 563

Configuring an ODBC data source for a DB2 data warehouse	564
Configuring Tivoli Performance Analyzer on Windows (ODBC connection)	565
Configuring Tivoli Performance Analyzer on Linux or UNIX (JDBC connection)	566
Starting Tivoli Performance Analyzer	568

Chapter 24. Tivoli Data Warehouse solution using DB2 on z/OS 569

Supported components	569
Prerequisite installation	571

Implementing a Tivoli Data Warehouse solution using DB2 on z/OS 571

Solution steps 573

Step 1: Connect the Warehouse Proxy node to your DB2 on z/OS database 573

Start defining the database connection	574
Define the communications protocol.	575
Define the TCP/IP communications parameters	576
Identify the DB2 on z/OS database	577
Register the database as an ODBC data source	578
Identify the z/OS server containing the DB2 on z/OS database	579
Define the DB2 on z/OS system options	580
Define the DB2 on z/OS security options	581

Complete the DB2 on z/OS host connection	582
Verify that the connection can be made.	583
Step 2: Configure the Tivoli Data Warehouse agents	585
Testing the database connection using the DB2	
Control Center	585
Testing the database connection using the DB2	
command-line processor.	588
Scenario: Connecting a Tivoli Data Warehouse	
solution using DB2 on z/OS to a Tivoli Enterprise	
Monitoring Server running z/OS.	589

Chapter 25. Tivoli Data Warehouse solution using Microsoft SQL Server . 595

Supported components	595
Prerequisite installation	597
Implementing a Tivoli Data Warehouse solution	
using Microsoft SQL Server.	598
Solution steps	598
Step 1: Create the Tivoli Data Warehouse database	599
Limiting the authority of the warehouse user	600
Step 2: Install and configure communications for	
the Warehouse Proxy Agent	600
Configuring an ODBC data source for a	
Microsoft SQL data warehouse	602
Configuring a Warehouse Proxy Agent on	
Windows (ODBC connection)	603
Configuring a Warehouse Proxy Agent on Linux	
or UNIX (JDBC connection)	605
Configuring the Warehouse Proxy Agent on	
Linux or UNIX: command-line procedure	608
Starting the Warehouse Proxy Agent.	609
Step 3: Configure communications between the	
Tivoli Enterprise Portal Server and the data	
warehouse	609
Configuring the portal server (ODBC	
connection)	610
Step 4: Install and configure communications for	
the Summarization and Pruning Agent.	612
Step 5: Install and configure communications for	
Tivoli Performance Analyzer	613
Configuring an ODBC data source for a	
Microsoft SQL data warehouse	614
Configuring Tivoli Performance Analyzer on	
Windows (ODBC connection)	615
Configuring Tivoli Performance Analyzer on	
Linux or UNIX (JDBC connection)	616
Starting Tivoli Performance Analyzer	617

Chapter 26. Tivoli Data Warehouse solution using Oracle 619

Supported components	619
Prerequisite installation	621
Implementing a Tivoli Data Warehouse solution	
using Oracle.	622
Solution steps	622
Step 1: Create the Tivoli Data Warehouse database	623
Creating the warehouse database on Oracle	623
Step 2: Install and configure communications for	
the Warehouse Proxy Agent	624
Creating a TNS Service Name	626

Configuring an ODBC data source for an Oracle	
data warehouse	627
Configuring a Warehouse Proxy Agent on	
Windows (ODBC connection)	627
Configuring a Warehouse Proxy Agent on Linux	
or UNIX (JDBC connection)	629
Starting the Warehouse Proxy	632
Step 3: Configure communications between the	
Tivoli Enterprise Portal Server and the data	
warehouse	632
Configuring a Windows portal server (ODBC	
connection)	633
Configuring a Linux or AIX portal server (JDBC	
connection)	634
Starting the portal server	636
Step 4: Install and configure communications for	
the Summarization and Pruning Agent.	636
Step 5: Install and configure communications for	
Tivoli Performance Analyzer	637
Creating a TNS Service Name	638
Configuring an ODBC data source for an Oracle	
data warehouse	639
Configuring a Tivoli Performance Analyzer on	
Windows (ODBC connection)	640
Configuring Tivoli Performance Analyzer on	
Linux or UNIX (JDBC connection)	641
Starting Tivoli Performance Analyzer	642

Chapter 27. Tivoli Data Warehouse solutions: common procedures . . . 643

Configuring the Summarization and Pruning Agent	
(JDBC connection).	643
Configuring the Warehouse Proxy Agent on Linux	
or UNIX: command-line procedure	653
Configuring the Summarization and Pruning Agent	
on Linux or UNIX: command-line procedure	654
Starting the Summarization and Pruning Agent	656
Installing and configuring multiple Warehouse	
Proxy Agents	656
Setting a permanent socket address for a proxy	
agent	657
Verifying the configuration.	657
Running the warehouse agents autonomously	658
Configuring a Warehouse Proxy Agent to run	
autonomously	658
Configuring a Summarization and Pruning	
Agent to run autonomously	659
Testing the connection between the portal server	
and the Tivoli Data Warehouse	663
1. Create the query	663
2. Create a workspace	664
3. Assign the query to the workspace	665
Tuning the performance of the Warehouse Proxy	666
Database initialization	667
Work queue	667
Connection pool	667
RPC threads and export requests.	668
Timeout values.	668
Configuring the Warehouse Proxy Agent and	
Summarization and Pruning Agent using the new	
graphical interface.	669

Installing and configuring the Warehouse Proxy Agent and Summarization and Pruning Agent in silent mode, using the silent response file	669
Installing and configuring the Warehouse Proxy Agent in silent mode on UNIX, using the silent response file.	670
Installing and configuring the Summarization and Pruning Agent in silent mode on UNIX, using the silent response file.	671
Upgrading the Warehouse Proxy Agent and Summarization and Pruning Agent in silent mode, using the silent response file	673
Upgrading the Warehouse Proxy Agent and Summarization and Pruning Agent using remote deploy.	673
Warehouse Proxy Agent silent install response file sample	674
Warehouse Proxy Agent silent config response file sample	674
Summarization and Pruning Agent silent install response file sample	677
Summarization and Pruning Agent silent config response file sample	677
Configuring the Warehouse Proxy Agent using the Tivoli Enterprise Portal	682
Configuring the Summarization and Pruning Agent using the Tivoli Enterprise Portal.	683
Remotely starting and stopping the Warehouse Proxy Agent using the Tivoli Enterprise Portal	684
Remotely starting and stopping the Summarization and Pruning Agent using the Tivoli Enterprise Portal	685
Remotely deploying the Warehouse Proxy Agent	685
Remotely deploying the Summarization and Pruning Agent	687
WAREHOUSELOG and	
WAREHOUSEAGGREGLOG tables	689

Chapter 28. Using Tivoli Common Reporting and monitoring agent reports 693

Chapter 29. Integrating event management systems. 695

Chapter 30. Setting up event forwarding to Netcool/OMNIBus . . . 697

Architecture overview	697
Event behavior	700
IBM Tivoli Monitoring Event Synchronization component	714
Architecture scenarios	715
One hub Tivoli Enterprise Monitoring Server and one Netcool/OMNIBus ObjectServer	715
Multiple hub Tivoli Enterprise monitoring servers and one Netcool/OMNIBus ObjectServer	718
One hub Tivoli Enterprise Monitoring Server and multiple Netcool/OMNIBus ObjectServers	721

Multiple IBM Tivoli Monitoring agents and one Netcool/OMNIBus ObjectServer	724
Netcool/OMNIBus Multitiered and High-availability Architecture	728
Integration with Tivoli Business Service Manager	735
Installation and configuration	738
Installing the IBM Tivoli Monitoring Event Synchronization Component	739
Installing from a wizard.	740
Installing from the command-line	743
Installing from the command-line using a silent installation	746
Updating the Netcool/OMNIBus ObjectServer with IBM Tivoli Monitoring attributes, tables, and triggers	749
Starting and stopping the IBM Tivoli Monitoring Situation Update Forwarder	759
Configuring the Netcool/OMNIBus EIF probe	760
Configuring the OMNIBus server for program execution from scripts	763
Configuring the hub monitoring server to forward events	764
Verifying installation and configuration.	765
Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIBus integration	766
Upgrading the IBM Tivoli Monitoring Event Synchronization Component	768
Upgrading existing IBM Tivoli Monitoring events to use new OMNIBus attributes	769
Replacing the default deduplication trigger	770
Updating the EIF Probe Rules	771
Customizing Event Integration	772
Configuring IBM Tivoli Monitoring Situation Update Forwarder event flow to OMNIBus	773
Changing the configuration values of the IBM Tivoli Monitoring Situation Update Forwarder	773
Updating the IBM Tivoli Monitoring Situation Forwarder to forward event status updates to additional monitoring servers	774
Customizing how the IBM Tivoli Monitoring OMNIBus triggers handle event status updates from the monitoring servers	775
Changing the default acknowledgment timeout used when sampled events are deleted or cleared in Netcool/OMNIBus	778
Editing the configuration parameters that the hub monitoring server uses to forward events	778
Specifying which situations forward events to Netcool/OMNIBus	780
Converting from unidirectional architecture to bidirectional architecture	781
Customizing event status processing behavior when agent switching is used or the agent goes offline	782
Creating a Netcool/OMNIBus WebGUI tool to launch from WebGUI to the Tivoli Enterprise Portal	785
Understanding and Customizing the Event Contents	787

Default mapping of situation events to OMNIBus events	787	Linux or UNIX Tivoli Enterprise Monitoring Automation Server worksheet	835
Generic mapping for agent-specific slots	792	Windows remote monitoring server worksheet	835
Adding event classes for new and updated agents to the MCS Attribute service used by the EIF Slot Customization function	793	Linux or UNIX remote monitoring server installation worksheet	836
About customizing the format of the event message attribute	795	Windows portal server worksheet	836
Localizing the event message attribute	796	Linux portal server worksheet	837
		Generic Windows monitoring agent worksheet	838
		Generic Linux or UNIX monitoring agent worksheet	838
		Windows portal desktop client worksheet	838
		Linux portal desktop client worksheet	839
		Monitoring server communications protocol details worksheet	839
Chapter 31. Setting up event forwarding to Tivoli Enterprise Console.	799		
Event integration with Tivoli Enterprise Console	800	Appendix B. Performing a silent installation of IBM Tivoli Monitoring	841
One or more hub monitoring servers and a single event server	800	Creating and using a Windows response file	841
A single hub monitoring server and multiple event servers	801	Automatically creating agent response files on Windows	842
Multiple hub monitoring servers and multiple event servers in a hub and spoke configuration	802	Running the silent installation from the command-line with parameters	844
Determining when to use the IBM Tivoli Enterprise Console	803	Running the silent installation using SMS	845
Installing event synchronization on your event server	804	Installing and configuring the Tivoli Enterprise Monitoring Automation Server in silent mode, by using the silent response file	845
Installing from a wizard	805	Installing and configuring the dashboard data provider in silent mode, by using the silent response file	847
Installing from the command-line	809	Performing a silent installation on a Linux or UNIX computer	848
Installing from the command-line using a silent installation	813	Installing components with a response file	849
Manually importing the event synchronization class files and rule set	816	Configuring components with a response file	851
Installing monitoring agent .baroc files on the event server	819	Automatically creating agent response files on Linux or UNIX	852
Configuring your monitoring server to forward events	820	Installing and configuring the Tivoli Enterprise Monitoring Automation Server in silent mode, by using the silent response file	853
Controlling event forwarding	821	Installing and configuring the dashboard data provider in silent mode, by using the silent response file	854
Starting and stopping the Situation Update Forwarder process	821	Performing a silent installation of software packages using IBM Installation Manager	855
Changing the configuration of the event synchronization component on the event server	822	Installing or updating software packages in silent mode	855
Defining additional monitoring servers to the event server	822	Installing or updating silently from authenticated repositories with IBM Installation Manager	858
Customizing event status processing behavior when agent switching is used or the agent goes offline	823	Recording a response file	858
Changing the TCP/IP timeout setting on your event server	825	Agent data location	859
Upgrading to Tivoli Event Synchronization version 3.0.0.0	826		
Upgrading from a wizard	826	Appendix C. Firewalls.	861
Upgrading from the command-line	827	Determining which option to use	861
Upgrading from the command-line using a silent installation	829	Flow of connection establishment	861
		Permission at the firewall	861
		Server address continuity	862
		Number of internet zones	862
		Basic (automatic) implementation	862
		Implementation with ephemeral pipe	863
		Implementation with partition files	864
Appendix A. Installation worksheets	833		
Windows hub monitoring server worksheet	833		
Windows Tivoli Enterprise Monitoring Automation Server worksheet	834		
Linux or UNIX hub monitoring server installation worksheet	834		

Sample scenarios	865
Windows: Editing the partition file	866
UNIX and Linux: Editing the partition file.	866
Creating the partition file manually	867
Implementation with firewall gateway	868
Configuration	869
Warehouse Proxy Configuration	872
Example gateway configuration scenario	873
Firewall network address translation (NAT) or multiple network interface cards	876
Defining a Tivoli Enterprise Portal Server interface on Windows	876
Defining a Tivoli Enterprise Portal Server interface on Linux or UNIX	877
Firewall scenarios for Tivoli Enterprise Portal	878

Appendix D. IBM Tivoli product, platform, and component codes . . . 883

Appendix E. Agent configuration and environment variables 887

Configuration files preserved during an upgrade	887
Preserving user custom settings	887
Files that are preserved on upgrade	888
Files that are not preserved on upgrade	889
Product behavior with custom configuration settings	890
General operations	890
Persistent configuration changes	891
Windows persistent configuration changes	891
Linux or UNIX persistent configuration changes	891
Environment variables	893
Common environment variables	893
Tivoli Enterprise Portal Server environment variables	894
Tivoli Enterprise Monitoring Server environment variables	899
Tivoli Enterprise Monitoring Automation Server environment variables	906
Tivoli Data Warehouse environment variables	907
Tivoli Monitoring Service Console variables	910
Common agent environment variables	911
Operating system agent environment variables	918

Appendix F. Maintaining the EIB on Linux or UNIX 923

Appendix G. Securing your IBM Tivoli Monitoring installation on Linux or UNIX 925

Usage	925
Scenario with secureMain	926
Securing your IBM Tivoli Monitoring environment at installation	926

Appendix H. Uninstalling IBM Tivoli Monitoring 929

Uninstalling the entire IBM Tivoli Monitoring environment.	929
--------------------------------------------------------------------	-----

Uninstalling the environment on Windows	929
Uninstalling the environment on Linux or UNIX	933
Uninstalling an individual IBM Tivoli Monitoring agent or component	934
Uninstalling a component on Windows.	934
Uninstalling a component on Linux or UNIX	935
Uninstalling OMEGAMON Monitoring Agents	935
Removing an agent through the Tivoli Enterprise Portal	937
Uninstalling the Warehouse Proxy	937
Removing the ODBC data source connection	938
Uninstalling components and agents silently	938
Performing a silent uninstallation on a Windows computer.	938
Performing a silent uninstallation on a Linux or UNIX computer	939
Uninstalling the event synchronization component	940
Uninstalling event synchronization manually	942
Uninstalling IBM Tivoli Monitoring components for a dashboard environment	944
Uninstalling IBM Tivoli Monitoring components for a dashboard environment in console mode	944
Performing a silent uninstallation of IBM Tivoli Monitoring components for a dashboard environment.	945
Silent response file for uninstallation of IBM Tivoli Monitoring components for a dashboard environment.	946
Uninstalling the Tivoli Authorization Policy Server or tivcmd Command Line Interface for Authorization Policy using the graphical user interface	947
Uninstalling the IBM Infrastructure Management Dashboards for Servers using the graphical user interface	947

Appendix I. Additional resources . . . 949

IBM Tivoli Monitoring 6 Welcome Kit	949
General education and support Web sites	949
Product documentation and IBM Redbooks	950
Education offerings	950
Service offerings	951
Other resources.	951

Appendix J. Documentation library 953

IBM Tivoli Monitoring library	953
Documentation for the base agents	954
Related publications	955
Other sources of documentation	955

Appendix K. Support for problem solving 957

Using IBM Support Assistant	957
Installing the IBM Support Assistant add-on	957
Collection of diagnostic files for self-diagnosis only	958
Obtaining fixes	958
Receiving weekly support updates	958
Contacting IBM Software Support	959
Determining the business impact of a problem	960

Describing problems and gathering information	960	Privacy policy considerations	965
Submitting problems	960	Index	967
Notices	963		
Trademarks	965		

Figures

1. IBM Tivoli Monitoring environment	7
2. IBM Tivoli Monitoring environment using Open Services Lifecycle Collaboration for product integration	9
3. Event synchronization overview	14
4. Tivoli Monitoring V6.3 communications model	22
5. Tivoli Monitoring component architecture including firewall gateway	23
6. Multiple data center environment	27
7. Warehouse load projection spreadsheet	38
8. Architecture of agentless monitoring	58
9. Adding agentless monitors to the deployment depot	61
10. Multiple domains with a common set of authorization policies	65
11. Tivoli Monitoring Startup Center Welcome page	106
12. Tivoli Monitoring Startup Center Installation Type window: Select the installation type	107
13. Tivoli Monitoring Startup Center Installation Type window	109
14. Tivoli Monitoring Startup Center Discover Machines window	110
15. Tivoli Monitoring Startup Center Assign Server Machines window	111
16. Tivoli Monitoring Startup Center Post Installation Summary window	113
17. Configuration window for the portal server database using DB2 for Linux, UNIX, and Windows	174
18. Configuration window for the Tivoli Data Warehouse database using DB2 for Linux, UNIX, and Windows	176
19. Configuration window for the Tivoli Data Warehouse database using Microsoft SQL Server	176
20. Configure Tivoli Performance Analyzer window - ODBC	178
21. Configure Tivoli Performance Analyzer window - JDBC	179
22. Configure Tivoli Performance Analyzer window - Advanced configuration	180
23. Configure Tivoli Performance Analyzer SPSS Configuration	181
24. Configure Tivoli Performance Analyzer Enable SPSS Configuration	182
25. Configure Tivoli Performance Analyzer	182
26. Sources pane of Configure Summarization and Pruning Agent window	183
27. Scheduling pane of Configure Summarization and Pruning Agent window	185
28. Work Days pane of Summarization and Pruning Agent configuration window	186
29. Log Settings pane of Summarization and Pruning Agent configuration window	187
30. Additional Settings pane of Summarization and Pruning Agent configuration window	188
31. Manage Tivoli Enterprise Monitoring Services window	190
32. Progress bar for application seeding	200
33. Install Product Support window	211
34. Configure Tivoli Enterprise Monitoring Automation Server.	221
35. Configure Tivoli Enterprise Monitoring Automation Server.	226
36. The Select Database for Tivoli Enterprise Portal window	230
37. Communications protocol between the portal server and the hub monitoring server	230
38. Tivoli Enterprise Portal Server Configuration	232
39. LDAP configuration	233
40. Authorization Policy Server Configuration	234
41. Reconfigure warehouse connection information for the portal server	235
42. Warehouse Proxy Database Selection	236
43. Configuration window for the portal server database using DB2 for Linux, UNIX, and Windows	237
44. Common Event Console Configuration window	239
45. Common Event Console Configuration window	252
46. Registering the portal server with the Tivoli Enterprise Monitoring Server	253
47. Configuring database connections for the portal server	254
48. Configuration information for the Tivoli Data Warehouse using an Oracle database.	256
49. Configure Dashboard data provider	258
50. LDAP Configuration	260
51. Installing the Agent Compatibility Package (component code AC).	266
52. Java Runtime Environment Not Detected error	267
53. The Select the Application Support to Add to the TEMS window	283
54. Application Support Addition Complete window	284
55. Refresh Configuration menu option	286
56. Manage Tivoli Enterprise Monitoring Services Install Product Support window	297
57. Java Control Panel window	346
58. Launching Java Web Start client	347
59. Java Plug-in	349
60. IBM Java 7 installation default language	349
61. IBM Java 7 Welcome page	350
62. IBM Java 7 Software License Agreement window	350
63. IBM Java 7 Choose Destination Location	351
64. Install System JVM.	351
65. Overwrite previous System JVM	351

66.	IBM Java 7 Start Copying Files.	352	106.	Tivoli Data Warehouse solution using DB2 on z/OS	570
67.	IBM Java 7 Browser Selection	352	107.	DB2 Client Configuration Assistant screen	574
68.	IBM Java 7 InstallShield Wizard Complete	353	108.	DB2 Add Database Wizard notebook, Source tab	575
69.	Upgrading Java applet cache	353	109.	DB2 Add Database Wizard notebook, Protocol tab	576
70.	IBM Java 5 application icon.	355	110.	DB2 Add Database Wizard notebook, TCP/IP tab	577
71.	Java Control Panel.	355	111.	DB2 Add Database Wizard notebook, Database tab	578
72.	Java Runtime Environment list.	356	112.	DB2 Add Database Wizard notebook, Data Source tab	579
73.	Java Control Panel Advanced tab	356	113.	DB2 Add Database Wizard notebook, Node Options tab	580
74.	Java Plug-in is not installed.	357	114.	DB2 Add Database Wizard notebook, System Options tab	581
75.	Download .rpm package file	358	115.	DB2 Add Database Wizard notebook, Security Options tab	582
76.	IBM Java 7 installing packages.	358	116.	DB2 Add Database Wizard notebook, DCS Options tab	583
77.	IBM Java 7 successfully installed	358	117.	Connection-confirmation screen	584
78.	Upgrading Java applet cache	360	118.	Connect to DB2 Database screen	584
79.	Java Control Panel	362	119.	DB2 Connection Confirmation screen	585
80.	Java Runtime Environment Settings	362	120.	DB2 Control Center	586
81.	Java Control Panel.	363	121.	DB2 Control Center right-click action menu	587
82.	Deployment Status Summary workspace showing the status of SSM deployments	382	122.	Connect to named database window	588
83.	Bulk deployment processing model	383	123.	DB2 command-line Processor window	589
84.	Output of the cinfo command	405	124.	Warehouse Proxy Database Selection screen	590
85.	Restart Component window: Tivoli Enterprise Monitoring Server	413	125.	Configure DB2 Data Source for Warehouse Proxy window	591
86.	Configure Tivoli Enterprise Monitoring Automation Server.	415	126.	Sources pane of Configure Summarization and Pruning Agent window	593
87.	Configure Tivoli Enterprise Monitoring Automation Server.	417	127.	Tivoli Data Warehouse solution using Microsoft SQL Server	596
88.	Restart of Monitoring Agent window	419	128.	Configure SQL Data Source for Warehouse Proxy window	604
89.	Hierarchy for the heartbeat interval	426	129.	Configure Warehouse Proxy window (Database Type).	606
90.	Manage Tivoli Enterprise Monitoring Services Advanced Utilities window	430	130.	Configure Warehouse Proxy window (Agent Parameters tab)	607
91.	The Manage Tivoli Enterprise Monitoring Services select the new portal server database window	431	131.	Configure SQL Data Source for Warehouse window	611
92.	The Manage Tivoli Enterprise Monitoring Services select the new portal server database window	431	132.	Tivoli Data Warehouse solution using Oracle	620
93.	Tivoli Enterprise Portal Server snapshot request screen	453	133.	Configure Oracle Data Source for Warehouse Proxy window	628
94.	Tivoli Enterprise Portal Server snapshot verification screen	454	134.	Configure Warehouse Proxy window (Database Type pane)	630
95.	Topology update buffering algorithm example	477	135.	Configure Warehouse Proxy window (Agent Parameters pane)	631
96.	Summary of support for the Tivoli Data Warehouse	520	136.	Configure Oracle Data Source for Warehouse window	633
97.	Configure Warehouse Proxy - Database Table Partitioning	532	137.	Configuring the connection to an Oracle data warehouse	635
98.	Configure Summarization and Pruning Agent - Database Table Partitioning	534	138.	Sources pane of Configure Warehouse Summarization and Pruning Agent window	646
99.	Tivoli Data Warehouse solution using DB2 for Linux, UNIX and Windows.	540	139.	Scheduling pane of Configure Warehouse Summarization and Pruning Agent window	648
100.	Warehouse Proxy Database Selection screen	553	140.	Work Days pane of Configure Warehouse Summarization and Pruning Agent window	649
101.	Configure DB2 Data Source for Warehouse Proxy window	554			
102.	Configure Warehouse Proxy window (Database Type).	556			
103.	Configure Warehouse Proxy window (Agent Parameters)	557			
104.	Configure DB2 Data Source for Warehouse window	560			
105.	Configuring the connection to a DB2 for Linux, UNIX and Windows data warehouse	561			

141. Log Settings pane of Configure Warehouse Summarization and Pruning Agent window	651	156. Single hub monitoring server and multiple event servers.	802
142. Additional Settings pane of Configure Warehouse Summarization and Pruning Agent window	652	157. Multiple hub monitoring servers and multiple event servers in a hub and spoke configuration	803
143. Create Query window	664	158. Window shown when no Tivoli Enterprise Console event server is found.	806
144. Assigning the WAREHOUSELOG query to a workspace	666	159. Upgrade data collection window	827
145. A typical event flow showing both IBM Tivoli Monitoring and Netcool/OMNIBus environments for a single-tier architecture	699	160. The Generate Response Files option	843
146. One hub Tivoli Enterprise Monitoring Server and one Netcool/OMNIBus ObjectServer	716	161. Structure of firewall gateway XML configuration document	870
147. Multiple hub Tivoli Enterprise Monitoring Servers and one Netcool/OMNIBus ObjectServer	719	162. Three-hop firewall scenario	873
148. One hub Tivoli Enterprise Monitoring Server and multiple Netcool/OMNIBus ObjectServer	722	163. Intranet with integral Web server	878
149. Multiple IBM Tivoli Monitoring Agents and one Netcool/OMNIBus ObjectServer	725	164. Intranet with external Web server.	879
150. Standard multitiered architecture with high availability	730	165. Intranet with integral Web server; Internet with external Web server.	880
151. single-tier high-availability architecture	731	166. Intranet and Internet with integral and external Web servers	881
152. Integrating IBM Tivoli Monitoring, Netcool/OMNIBus, and Tivoli Business Service Manager	736	167. Two host addresses, intranet and Internet, with integral and external Web servers	882
153. Installation of IBM Tivoli Monitoring and Tivoli Event Synchronization	740	168. Uninstalling IBM Tivoli Monitoring	930
154. Netcool/OMNIBus event synchronization configuration fields	741	169. Confirming the uninstallation	930
155. One or more hub monitoring servers connecting to a single event server	801	170. Stopping Tivoli components before uninstallation	931
		171. Removing the portal database	931
		172. Database information	931
		173. Uninstallation progress window	932
		174. GSKit uninstallation	932
		175. Successful uninstallation	933

Tables

1. Planning checklist	19
2. Warehouse database server considerations	39
3. Portal client deployment considerations	40
4. Example distributed installation.	41
5. Support files to be installed	41
6. Data collectors usable with the various agentless monitors and releases supported . .	59
7. User's guides for the agentless monitors	61
8. Update history for the baroc files for IBM Tivoli Monitoring agents and components . .	68
9. Staffing estimates	72
10. Supported Windows operating systems	75
11. Supported UNIX, IBM i, and z/OS operating systems	77
12. Supported Linux operating systems	79
13. Operating system requirements for IBM GSKit	81
14. Supported databases for the portal server	82
15. Supported databases for the Tivoli Data Warehouse	83
16. Supported databases for Tivoli Data Warehouse database compression	84
17. Supported databases for Tivoli Data Warehouse database range partitioning . . .	84
18. Estimated memory and disk storage for IBM Tivoli Monitoring components on distributed systems	85
19. Required software for IBM Tivoli Monitoring	89
20. Supported browser versions	90
21. Commands for determining your system's kernel version.	97
22. Startup Center: Supported Windows operating systems	104
23. Startup Center: Supported Linux operating systems	104
24. NetBIOS Reserved Ports	116
25. Installation and configuration steps	125
26. Supported IBM Tivoli Monitoring configurations using the IPv6 communications protocol	128
27. Agents requiring warehouse database migration	151
28. Upgrading from IBM Tivoli Monitoring V6.1 or V6.2 to IBM Tivoli Monitoring V6.3 . . .	161
29. Configuration information for the portal server database	174
30. Configuration information for the Tivoli Data Warehouse database	176
31. Tivoli Data Warehouse URLs	184
32. JDBC driver names	184
33. IBM Tivoli Monitoring high-level installation steps	194
34. Communications protocol settings for the hub monitoring server	198
35. Steps for installing a hub monitoring server on a Linux or UNIX computer	200
36. UNIX monitoring server protocols and values	203
37. tacmd commands to control the operational state of the hub self-describing agent function	205
38. Remote monitoring server communications protocol settings	214
39. Steps for installing a remote monitoring server on a Linux or UNIX computer . . .	215
40. UNIX monitoring server protocols and values	216
41. Configuration information for the Authorization Policy Server.	234
42. Configuration information for the portal server database	237
43. Steps for installing a portal server on a Linux or AIX computer	241
44. Hub monitoring server protocols and values	245
45. Configuration information for the Authorization Policy Server.	249
46. Configuration information for the Tivoli Enterprise Portal Server database	255
47. Configuration information for the Authorization Policy Server.	259
48. Communications protocol settings	264
49. Steps for installing a monitoring agent on Linux or UNIX	267
50. UNIX monitoring server protocols and values	269
51. Procedures for installing and enabling application support	277
52. Product support on the Infrastructure and Agent DVDs.	278
53. Tivoli Performance Analyzer Reports and Domain Definitions	278
54. Installation media and instructions for installing application support for nonbase monitoring agents	280
55. Locations of CAT and ATR files for the monitoring server	292
56. Locations of application support files on a Linux or UNIX monitoring server.	294
57. Locations of CAT and ATR files for the monitoring server	295
58. Installation, update, and configuration topics	307
59. Installation and configuration topics	319
60. Tivoli Enterprise Portal client support	344
61. File locations for changing application properties for UNIX and Linux	364
62. Remote agent deployment tasks	371
63. Agent depot management commands	374
64. Interaction between member properties and group properties	388
65. Property precedence between deploy groups and bundle groups.	389
66. tacmd commands to manage your self-describing agent installations	394
67. Configuration tasks available through Manage Tivoli Enterprise Monitoring Services . . .	411
68. Communications protocol settings	412
69. Communications protocol settings	418

70. Using COUNT and SKIP variables to assign port numbers	425	98. Configuration information for the Tivoli Data Warehouse database on DB2 running on z/OS	591
71. Attribute groups that must be configured	433	99. Goals for achieving a Tivoli Data Warehouse solution using Microsoft SQL Server	595
72. Browser client behavior	458	100. Information topics related to installation of prerequisite software for a Tivoli Data Warehouse solution	598
73. Overview of SOAP Server configuration steps	461	101. Default values for Tivoli Data Warehouse parameters	599
74. TCP/IP Fields in Hub Specification dialog	462	102. Tasks for installing and configuring communications for the Warehouse Proxy Agent	601
75. SNA Fields in Hub Specification dialog	462	103. Configuration information for the Tivoli Data Warehouse database on Microsoft SQL Server	604
76. SOAP hub configuration values	464	104. Tasks for configuring communications between the portal server and a Microsoft SQL Server data warehouse	609
77. SOAP methods associated with access privileges	464	105. Configuration information for the Tivoli Data Warehouse database on Microsoft SQL Server	611
78. Default Java heap size parameters by portal client type	478	106. Tasks for installing and configuring communications for the Summarization and Pruning Agent	612
79. Tivoli Data Warehouse database size estimation worksheet	513	107. Tasks for installing and configuring communications for the Tivoli Performance Analyzer	613
80. Database size examples	514	108. Configuration information for the Tivoli Data Warehouse database on Microsoft SQL	616
81. Goals for creating a Tivoli Data Warehouse solution using DB2 for Linux, UNIX, and Windows	539	109. Goals for achieving a Tivoli Data Warehouse solution using Oracle	619
82. Information topics related to installation of prerequisite software for a Tivoli Data Warehouse solution	542	110. Information topics related to installation of prerequisite software for a Tivoli Data Warehouse solution	621
83. Tasks for creating the Tivoli Data Warehouse database	543	111. Tasks for creating the Tivoli Data Warehouse database	623
84. Default values for Tivoli Data Warehouse parameters	544	112. Default values for Tivoli Data Warehouse parameters	623
85. Tasks for installing and configuring communications for the Warehouse Proxy Agent	548	113. Tasks for installing and configuring communications for the Warehouse Proxy Agent	624
86. Configuration information for the Tivoli Data Warehouse database on DB2 for Linux, UNIX and Windows	554	114. Configuration information for the Tivoli Data Warehouse database on Oracle	628
87. Tasks for configuring communications between the portal server and a DB2 for Linux, UNIX and Windows data warehouse	559	115. Tasks for configuring communications between the portal server and an Oracle data warehouse	632
88. Configuration information for the Tivoli Data Warehouse database on DB2 for Linux, UNIX and Windows	560	116. Configuration information for the Tivoli Data Warehouse database on Oracle	634
89. Configuration information for the Tivoli Data Warehouse database on DB2 for Linux, UNIX and Windows	562	117. Configuration information for a Tivoli Data Warehouse database on Oracle	635
90. Tasks for installing and configuring communications for the Summarization and Pruning Agent	562	118. Tasks for installing and configuring communications for the Summarization and Pruning Agent	636
91. Tasks for installing and configuring communications for the Tivoli Performance Analyzer	564	119. Tasks for installing and configuring communications for the Tivoli Performance Analyzer	637
92. Configuration information for the Tivoli Data Warehouse database on DB2 for Linux, UNIX and Windows	566	120. Configuration information for the Tivoli Data Warehouse database on Oracle	641
93. Goals for creating a Tivoli Data Warehouse solution using DB2 on z/OS	569	121. Where to obtain the JDBC driver files for the Summarization and Pruning Agent	644
94. Information topics related to installation of prerequisite software for a Tivoli Data Warehouse solution	571	122. Tivoli Data Warehouse URLs	646
95. Characteristics of implicitly created DB2 on z/OS database	572	123. JDBC driver names	647
96. Maximum DB2 on z/OS page sizes	572		
97. Required parameters for accessing the DB2 on z/OS database	572		

124. Descriptions of the columns in the WAREHOUSESUMPRUNE control settings table	661	145. EIF attribute slots to OMNIBus attributes, not to be customized	791
125. Behavior of events originating from a hub monitoring server	700	146. Special characters for attribute groups and names in EIF events generated from forwarded situation events.	793
126. Behavior of events originating from IBM Tivoli Monitoring agents.	712	147. Tivoli Enterprise Console event synchronization installation and configuration steps	799
127. Installation and configuration: one hub Tivoli Enterprise Monitoring Server and one Netcool/OMNIBus ObjectServer	717	148. IBM Tivoli Enterprise Console event synchronization configuration fields	806
128. Installation and configuration: multiple hub Tivoli Enterprise Monitoring Servers and one Netcool/OMNIBus ObjectServer	720	149. IBM Tivoli Enterprise Console event synchronization configuration fields, continued.	807
129. Installation and configuration: one hub Tivoli Enterprise Monitoring Server and multiple Netcool/OMNIBus ObjectServers	723	150. Options for rule base modification	808
130. Installation and configuration for agents sending events to the Netcool/OMNIBus Probe for Tivoli EIF	726	151. IBM Tivoli Enterprise Console event synchronization configuration values	814
131. Installation and configuration if agents are sending events to the Netcool/OMNIBus Probe for SNMP	727	152. Variables to customize the behavior of event status processing when agent switching is used.	824
132. Installation and configuration: Netcool/OMNIBus Multitiered and High-Availability Architecture ObjectServers	732	153. Installation and configuration steps	841
133. Installation and configuration	737	154. Silent installation parameters for UNIX	850
134. Netcool/OMNIBus event synchronization configuration fields	741	155. Component product codes for infrastructure components and base monitoring agents	883
135. Netcool/OMNIBus event synchronization configuration fields, continued.	742	156. Platform codes required for remote agent deployment	884
136. Netcool/OMNIBus event synchronization configuration values	747	157. Application support codes	886
137. ObjectServer updates required for each architecture type	749	158. Information topics related to IBM Tivoli Monitoring agent configuration and environment variables	887
138. Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIBus integration.	767	159. Common environment variables	893
139. Additional configuration tasks after event integration installation	772	160. Tivoli Enterprise Portal Server environment variables	894
140. Supported Netcool/OMNIBus Probe for Tivoli EIF configuration parameters for the Event Integration Facility (EIF).	779	161. Tivoli Enterprise Monitoring Server environment variables	899
141. Variables to customize the behavior of event status processing when agent switching is used	783	162. Tivoli Enterprise Monitoring Automation Server environment variables	906
142. Common tasks for understanding and customizing the contents of an event.	787	163. Warehouse proxy environment variables	908
143. Base IBM Tivoli Monitoring EIF attribute slots.	788	164. Summarization and Pruning agent environment variables	909
144. Mapping of EIF attribute slots to OMNIBus attributes	790	165. Service Console environment variables	911
		166. Common agent environment variables	911
		167. Windows OS monitoring agent environment variables	919
		168. Linux OS monitoring agent environment variables	919
		169. UNIX OS monitoring agent environment variables	920
		170. EIB Files	923
		171. Candle OMEGAMON Release 04R1	936
		172. Candle OMEGAMON Release BIV110	936

Chapter 1. Introduction

This chapter reviews the new features and enhancements that affect installation or configuration, and describes the architecture of the IBM® Tivoli® Monitoring products. This information will help you to plan your deployment and prepare to install, upgrade, or configure the product's base components.

Tivoli Monitoring products use a set of service components (known collectively as Tivoli Management Services) that are shared by a number of other product suites, including IBM Tivoli OMEGAMON® XE monitoring products, IBM Tivoli Composite Application Manager products, System Automation for z/OS®, Web Access for Information Management, and others. The information in this section is also relevant to these products.

New in this release

Review the latest enhancements to Tivoli Monitoring components that are relevant to the *IBM Tivoli Monitoring Installation and Setup Guide*.

New in Version 6.3 Fix Pack 1

The following enhancements to IBM Tivoli Monitoring affect installation or configuration since the release of Version 6.3.

IBM Infrastructure Management Dashboards for Servers Version 6.3 Fix Pack 1

You can update the IBM Infrastructure Management Dashboards for Servers to Version 6.3 Fix Pack 1. See the following topics for more information about updating the IBM Infrastructure Management Dashboards for Servers:

- “Updating the IBM Infrastructure Management Dashboards for Servers using the graphical user interface” on page 315
- “Installing or updating the IBM Infrastructure Management Dashboards for Servers using console mode” on page 311
- “Installing or updating software packages in silent mode” on page 855

TEMS Security Compatibility Mode is enabled by default

To ease the transition to Take Action command security, the KMS_SECURITY_COMPATIBILITY_MODE environment variable is enabled by default. When you are ready to implement strict command security (that is each command must have an encrypted security token) the TEMS Security Compatibility Mode variable should be changed from Y to N. This will cause all agents connected to that Tivoli Enterprise Monitoring Server to check for a valid command security token before allowing a command to execute. It also requires clocks to be synchronized to Coordinated Universal Time (UTC) across the systems with monitoring servers, the portal server, and monitoring agents, and Integrated Cryptographic Service Facility to be installed and configured on z/OS systems with monitoring servers. See the following topics for more information about configuring TEMS Security Compatibility Mode and the KMS_SECURITY_COMPATIBILITY_MODE environment variable:

- “Installing and configuring the hub Tivoli Enterprise Monitoring Server” on page 196
- “Tivoli Enterprise Monitoring Server environment variables” on page 899

New in Version 6.3

The following sections describe changes in this release that affect installation or configuration.

For a complete list of new product features, see the *IBM Tivoli Monitoring Administrator's Guide*.

Jazz™ for Service Management

Jazz for Service Management brings together the Open Services for Lifecycle Collaboration (OSLC) community's open specifications for linking data, shared administrative services, dashboard and reporting services. Through these facets, Jazz for Service Management accelerates deployment, integration, and workflow automation across IBM, partner, and third-party tools.

Jazz for Service Management has a number of integration services: Administration, Registry, IBM Tivoli Common Reporting, Security, and IBM Dashboard Application Services Hub. These integration services provide key features including:

- Shared data repository for products integrating through Jazz for Service Management.
- Consistent UI experience through Dashboard Application Services Hub in Jazz for Service Management.
- Simplified administration of products and solutions integrating through Jazz for Service Management.
- Ad hoc, self-service reporting through Tivoli Common Reporting in Jazz for Service Management.

For more information about Jazz for Service Management, go to the Jazz for Service Management Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/psc_ic-homepage.html).

IBM Infrastructure Management Dashboards for Servers running on the Dashboard Application Services Hub V3.1 or later

With the IBM Tivoli Monitoring dashboard data provider enabled, Dashboard Application Services Hub users can retrieve managed system groups and events for all monitoring agents and agent health metrics for Linux OS agent, UNIX OS agent, and Windows OS agent using the Infrastructure Management Dashboards for Servers application. This application is installed and configured into Dashboard Application Services Hub V3.1 or later using IBM Installation Manager. For more information, see “Installing and configuring the IBM Infrastructure Management Dashboards for Servers” on page 306 and “Required software and memory requirements for a dashboard environment” on page 91.

Dashboard data provider configuration option for the Tivoli Enterprise Portal Server

The IBM Tivoli Monitoring dashboard data provider retrieves monitoring agent data for display in the Dashboard Application Services Hub. The dashboard data provider is optionally enabled during the Tivoli Enterprise Portal Server configuration. With the dashboard data provider enabled and configured, Dashboard Application Services Hub users can retrieve read-only data from the hub monitoring server and monitoring agents for display in monitoring dashboards. Examples of such dashboards are the IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM®, or custom dashboards.

Dashboard data provider is the new name for the ITM REST service. In IBM Tivoli Monitoring V6.3, the dashboard data provider has been enhanced to improve performance for monitoring dashboards. For more information, see “Creating a connection to the IBM Tivoli Monitoring dashboard data provider” in the *IBM Tivoli Monitoring Administrator's Guide*.

Tivoli Authorization Policy Server

The Tivoli Authorization Policy Server feature provides you with greater access control capabilities. You can protect your resources from unauthorized access by users of monitoring dashboards in the IBM Dashboard Application Services Hub. IBM Tivoli Monitoring V6.3 with the Authorization Policy Server feature enabled provides the following capabilities:

- To restrict the access for dashboard users to Managed System Groups and to individual agent managed systems.
- The ability to grant role-based access control to users and user groups, making access control easier and safer.

To implement the feature, you must install IBM Installation Manager packages for the Tivoli Authorization Policy Server and the tivcmd Command Line Interface for Authorization Policy. The Authorization Policy Server is installed with your Dashboard Application Services Hub. The command-line interface is installed on the computers that are used by the administrators who are creating authorization policies. After successful installation of these two packages, you can execute various CLI commands as required to create roles, grant permissions, exclude permissions, and so on. After you have created roles and permissions, you enable authorization policy enforcement in the Tivoli Enterprise Portal Server configuration. For more information, see “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318 and “Required software and memory requirements for a dashboard environment” on page 91.

Installation Launchpad and IBM Installation Manager

IBM Installation Manager is a tool that is designed to install and maintain your software packages. You can use IBM Installation Manager to install the following IBM Tivoli Monitoring V6.3 components:

- Tivoli Authorization Policy Server
- tivcmd Command Line Interface for Authorization Policy
- IBM Infrastructure Management Dashboards for Servers

Wizards guide you through the steps to install, modify, update, roll back, or uninstall your IBM products. Use Installation Manager to install software packages on your local machine, or with the IBM Packaging Utility to install software for an enterprise from a network accessible repository. Installation Manager is installed along with the new IBM Tivoli Monitoring components, if it is not already installed on your machine.

For more information about IBM Installation Manager and the IBM Packaging Utility, go to <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.

If you want to install the Authorization Policy Server, tivcmd CLI, or IBM Infrastructure Management Dashboards for Servers from the installation media, you can use the IBM Tivoli Monitoring installation launchpad application. The installation launchpad automatically installs, configures, and starts IBM Installation Manager. Once started, you can use IBM Installation Manager to install and configure the components onto the local system. This approach assumes that you have already installed the IBM Dashboard Application Services Hub component and prerequisites by using one of the Jazz for Service Management installation technologies.

Alternatively, you can use the IBM Packaging Utility for IBM Installation Manager to create a network repository that includes the IBM Infrastructure Management Dashboards for Servers component, the Authorization Policy Server, and tivcmd CLI, as well as Jazz for Service Management components and prerequisites. On each system where you want to install an IBM Tivoli Monitoring component for a dashboard environment, you can configure IBM Installation Manager with the location of the network repository and then select the components to install on that system.

For more information about installing IBM Tivoli Monitoring components with IBM Installation Manager, see the following topics:

- “Required software and memory requirements for a dashboard environment” on page 91
- “Using IBM Installation Manager” on page 42
- “Using IBM Installation Manager to install or update IBM Tivoli Monitoring components for a dashboard environment” on page 299
- “Using the Installation Launchpad to install or update IBM Tivoli Monitoring components for a dashboard environment” on page 299
- For the Jazz for Service Management Installation Guide, installation scenarios, and technologies that are supported for the Jazz for Service Management components, go to the Jazz for Service Management Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/psc_ic-homepage.html).

Open Services Lifecycle Collaboration Service Provider

The Tivoli Enterprise Monitoring Automation Server component contains the Open Services Lifecycle Collaboration Performance Monitoring (OSLC-PM) service provider. The Tivoli Enterprise Monitoring Automation Server is installed on the same systems as your Hub Tivoli Enterprise Monitoring Server. The service provider registers monitoring resources with the Registry Services. Registry Services is a Jazz for Service Management integration service that provides a shared data repository for products in an integrated service management environment. Products that discover and manage shared IT resources can register these IT resources and the services they offer with Registry Services. Other products can consume data by querying Registry Services for the managed resources or the associated service providers of interest. In the query responses, Registry Services returns links to the service provider that can be queried to retrieve more information about the managed resources.

For more details about Performance Monitoring service provider and the Tivoli Enterprise Monitoring Automation Server, see the following topics:

- “Required software and memory requirements for using the OSLC Performance Monitoring service provider” on page 94
- “Tivoli Enterprise Monitoring Automation Server” on page 17
- “Installing and configuring Tivoli Enterprise Monitoring Automation Server” on page 219
- *Performance Monitoring service provider* in the *IBM Tivoli Monitoring Administrator's Guide*.

Native 64-bit versions of IBM Tivoli Monitoring components on Windows systems

You can install and configure 32-bit components from a separate installation image (intended for 32-bit Windows systems) or install and configure 64-bit server components from a separate installation image (intended for 64-bit Windows systems). In addition, the Tivoli Enterprise Portal browser and Java™ Webstart clients now support operation under a 64-bit Windows browser, for example Internet Explorer 64-bit. For this configuration, you must install a 64-bit JVM supplied by IBM, or from Oracle. Full support for 32-bit browsers under Windows, for example Internet Explorer and Firefox, is also provided.

For more information and the conditions that apply to upgrading the Tivoli Enterprise Monitoring Server and Tivoli Enterprise Portal Server, see “Prerequisites for IBM Tivoli Monitoring V6.3” on page 146.

Native 64-bit binaries on Linux for AMD and Intel systems

The Tivoli Enterprise Monitoring Server, Warehouse Proxy Agent, and Summarization and Pruning Agent now provide native 64-bit binaries on Linux for AMD and Intel systems. The Tivoli Enterprise Portal now provides native 64-bit binaries on Linux for Intel systems. All three Tivoli Enterprise Portal deployment modes - desktop, browser, and Java Webstart all support native 64-bit operation under a 64-bit Java Runtime Environment.

Tivoli Enterprise Portal browser support for 32-bit and 64-bit browsers

The Tivoli Enterprise Portal browser client now supports both 32-bit and 64-bit Firefox and Internet Explorer browsers.

Current versions of Mozilla Firefox now supported by the Tivoli Enterprise Portal browser client

Firefox 3.5 and earlier releases of Firefox are no longer supported with IBM Tivoli Monitoring V6.3. If Firefox is your preferred browser to use with the Tivoli Enterprise Portal, then use the most recent Firefox Extended Support Release (ESR) available from <http://www.mozilla.org>. Upgrading to a later version of Firefox eliminates the post-installation configuration of the JVM that is required to successfully operate the Tivoli Enterprise Portal browser client. For more information, see “Browser plug-in support for Java applets” on page 344.

Tivoli Data Warehouse range partitioning

Range partitioning is a database data organization feature that can significantly improve pruning and query performance in large Tivoli Data Warehouse databases. For more information and the associated configuration parameters, see Chapter 21, “Tivoli Data Warehouse range partitioning,” on page 531 and Chapter 22, “Specifying where tables and indices are created,” on page 537.

Filter by table in Schema Publication Tool

You can now filter by table in any mode of the Schema Publication Tool. For more information, see “Generating SQL for data warehouse tables” on page 525.

Changes to default self-describing agent behavior and new tacmd commands

In V6.2.3 and V6.2.3 Fix Pack 1, when the self-describing agent capability is turned on, any available products and versions are automatically installed. As of V6.3, the hub monitoring server blocks all self-describing agent installation until you issue one of the new tacmd commands. You can also use new tacmd commands to enable or disable the state of the self-describing agent function at the hub monitoring server, without having to recycle any hub monitoring servers.

These features provide more control over what products and versions are installed on your monitoring server and portal server by the automatic self-describing agent process. For more information and other new tacmd commands, see “Managing your self-describing agent installations” on page 394 and “Dynamically controlling the hub monitoring server self-describing agent capability” on page 205.

SOAP security enhancements and duper process optimization

You can now enable security for CT_Email and CT_Export requests by using the SOAP_IS_SECURE environment variable on the monitoring server.

For more information, see *Enabling SOAP security* in the *IBM Tivoli Monitoring Administrator's Guide*.

Tivoli Universal Agent and UNIX Log Agent not included in V6.3

The Tivoli Universal Agent and the UNIX Log Agent are not included with IBM Tivoli Monitoring Version 6.3. If you already have any of these agents installed in your environment, you can still use them with IBM Tivoli Monitoring Version 6.3.

Tivoli Integrated Portal name change

The V3.1 release of Tivoli Integrated Portal is now referred to as the Dashboard Application Services Hub.

i5/OS™ agent name change

The i5/OS monitoring agent is now referred to as the IBM i monitoring agent.

Overview of IBM Tivoli Monitoring

This chapter describes the architecture of the IBM Tivoli Monitoring products and provides information to help you plan your deployment and prepare to install, upgrade, or configure the product's base components.

IBM Tivoli Monitoring products monitor the performance and availability of distributed operating systems and applications. These products are based on a set of common service components, referred to collectively as Tivoli Management Services. Tivoli Management Services components provide security, data transfer and storage, notification mechanisms, user interface presentation, and communication services in an agent-server-client architecture (Figure 1 on page 7). These services are shared by a number of other products, including IBM Tivoli XE mainframe monitoring products and IBM Tivoli Composite Application Manager products, as well as other IBM Tivoli Monitoring products such as Tivoli Monitoring for Applications, Tivoli Monitoring for Cluster Managers, Tivoli Monitoring for Databases, Tivoli Monitoring for Energy Management, Tivoli Monitoring for Messaging and Collaboration, Tivoli Monitoring for Messaging and Collaboration, and Tivoli Monitoring for Virtual Environments.

This book contains information on deploying, installing, and configuring the common services components and the monitoring agents that comprise the base IBM Tivoli Monitoring product on distributed systems. If you purchased a product other than IBM Tivoli Monitoring that uses Tivoli Management Services, use this book to install and configure the common components. Do not install the base operating system agents listed below unless you have also licensed IBM Tivoli Monitoring.

Linux OS

UNIX OS

Windows OS

IBM i OS

Agentless Monitoring for AIX®

Agentless Monitoring for HP-UX

Agentless Monitoring for Linux

Agentless Monitoring for Solaris

Agentless Monitoring for Windows

If you purchased additional IBM Tivoli Monitoring products, see their documentation for agent-specific installation and configuration information.

IBM Tivoli Monitoring V6.3 products are enabled for use with IBM License Metric Tool. Tivoli Management Services components and Tivoli Monitoring agents provide inventory signature files and usage definitions that allow License Metric Tool to report installed products and product usage by computer. License Metric Tool support is an optional capability that requires License Metric Tool version 7.5 or later.

Components of the monitoring architecture

This section describes the architecture of the IBM Tivoli Monitoring products and provides information to help you plan your deployment and prepare to install or upgrade the base components of the product.

Tivoli Monitoring products use a set of service components (known collectively as Tivoli Management Services) that are shared by a number of other product suites, including IBM Tivoli XE monitoring products, IBM Tivoli Composite Application Manager products, System Automation for z/OS, Web Access for Information Management, and others. The information in this section is also relevant to these products.

Tivoli Monitoring products, and other products that share Tivoli Management Services, participate in a server-client-agent architecture. Monitoring agents for various operating systems, subsystems, databases, and applications (known collectively as Tivoli Enterprise Monitoring Agents) collect data and send it to a Tivoli Enterprise Monitoring Server. Data is accessed from the monitoring server by Tivoli Enterprise Portal clients and by dashboard users of the Dashboard Application Services Hub. A Tivoli Enterprise Portal Server provides presentation and communication services for these clients. Several optional components such as an historical data warehouse extend the functionality of the framework. Figure 1 on page 7 shows the configuration of an IBM Tivoli Monitoring environment.

IBM Tivoli Monitoring also includes Jazz for Service Management which brings together the Open Services for Lifecycle Collaboration (OSLC) community's open specifications for linking data and other shared integration services, including administrative, dashboard, reporting, and security services. You can use IBM Tivoli Monitoring components to extend the Jazz for Service Management functionality for your monitoring environment as shown in Figure 1 on page 7 and Figure 2 on page 9.

Before deciding where to deploy the components of the Tivoli Monitoring product in your environment, you should understand the components of the product, the roles that they play, and what affects the load on these components.

Architecture Overview

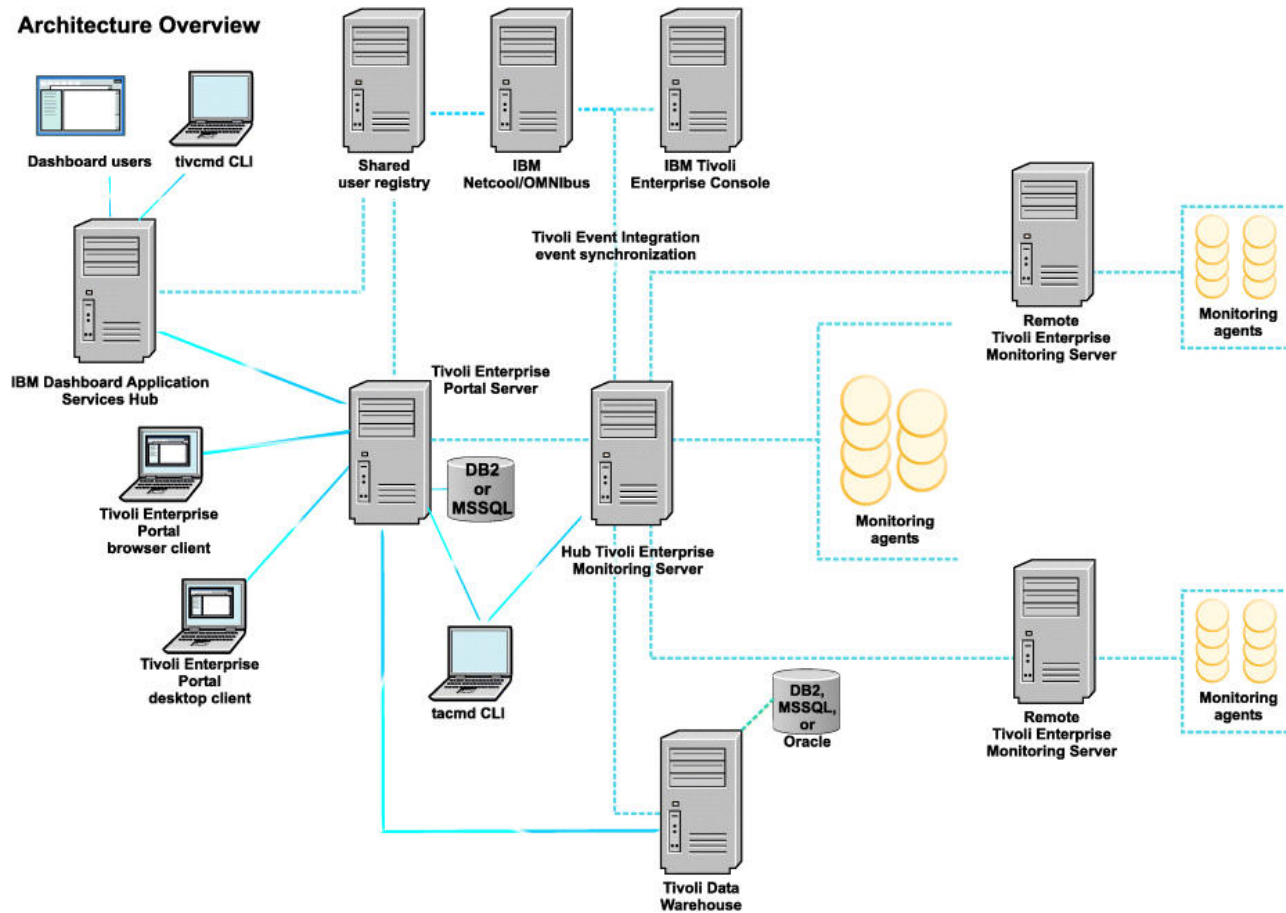


Figure 1. IBM Tivoli Monitoring environment

A typical environment comprises the following components:

- One or more Tivoli Enterprise Monitoring Servers, which act as a collection and control point for alerts received from the agents, and collect their performance and availability data. The monitoring server also manages the connection status of the agents. One server in each environment must be designated as the *hub*.
- A Tivoli Enterprise Portal Server, which provides the core presentation layer for retrieval, manipulation, analysis, and pre-formatting of data. The portal server retrieves data from the hub monitoring server in response to user actions at the portal client, and sends the data back to the portal client for presentation. The portal server also provides presentation information to the portal client so that it can render the user interface views suitably.
- One or more Tivoli Enterprise Portal clients, with a Java-based user interface for viewing and monitoring your enterprise. Tivoli Enterprise Portal offers two modes of operation: desktop and browser.
- Tivoli Enterprise Monitoring Agents, installed on the systems or subsystems you want to monitor. These agents collect data from monitored, or managed, systems and distribute this information either to a monitoring server or to an EIF or SNMP event server such as Netcool/OMNibus.
- One or more instances of the tacmd Command Line Interface (CLI). This CLI is used to manage your monitoring environment and can also be used to automate many of the administrative functions performed using the Tivoli Enterprise Portal. The CLI commands either send requests to the Hub monitoring server or to the Tivoli Enterprise Portal Server.

- z/OS only: Tivoli Management Services: Engine (TMS:Engine) provides common functions, such as communications, multithreaded runtime services, diagnosis (dumps), and logging (RKLVLLOG), for the Tivoli Enterprise Monitoring Server, monitoring agents, and components of XE products running on z/OS.
- An Eclipse Help Server for presenting help for the portal and all monitoring agents for which support has been installed. The help server is installed with Tivoli Enterprise Portal Server.

An installation optionally includes the following components:

- Tivoli Data Warehouse for storing historical data collected from agents in your environment. The data warehouse is located on an IBM DB2® for Linux, UNIX, and Windows, DB2 on z/OS, Oracle, or Microsoft SQL database. To store data in this database, you must install the Warehouse Proxy Agent. To perform aggregation and pruning functions on the data, you must also install the Summarization and Pruning Agent.
- Event synchronization component, the Event Integration Facility, that sends updates to situation events that have been forwarded to a Netcool/OMNIbus ObjectServer or a Tivoli Enterprise Console® event server back to the monitoring server.
- IBM Dashboard Application Services Hub is a Jazz for Service Management component that provides dashboard visualization and reporting services. Operators of the dashboard access it through a web browser interface. You can install the following types of applications into the Dashboard Application Services Hub:
 - The IBM Infrastructure Management Dashboards for Servers application displays situation event information, managed system groups and key health metrics for Windows OS agents, Linux OS agents, and UNIX OS agents. Situation events and monitoring data are retrieved from the Tivoli Enterprise Portal Server using its dashboard data provider.

Note:

1. Other monitoring products such as IBM Tivoli Monitoring for Virtual Environments and IBM Smart Cloud Monitoring might provide their own management dashboard applications that use the dashboard data provider. You can also create custom dashboard views that display monitoring data using the Dashboard Application Services Hub user interface.
 2. To use monitoring dashboards, you must enable the dashboard data provider component of the Tivoli Enterprise Portal Server. IBM Dashboard Application Services Hub sends requests for monitoring data to the dashboard data provider which uses the portal server services to retrieve agent data through the monitoring servers.
- The Tivoli Authorization Policy Server application is used to create authorization policies that control which managed system groups and managed systems can be viewed by a dashboard operator. The authorization policies are created using the tivcmd Command Line Interface for Authorization Policy and stored at the Authorization Policy Server. The policies are enforced in the dashboard data provider component of the Tivoli Enterprise Portal Server. The dashboard data provider retrieves the policies from the Authorization Policy Server.
 - Tivoli Common Reporting can be used to gather, analyze, and report important trends in your managed environment using historical data from the Tivoli Data Warehouse. The Tivoli Common Reporting user interface is installed with Dashboard Application Services Hub and can be used to display predefined reports provided by monitoring agents and to create custom reports. Tivoli Common Reporting accesses the Tivoli Data Warehouse directly (this interaction is not depicted in Figure 1 on page 7).
- A shared user registry is an LDAP server such as Tivoli Directory Server or Microsoft Active Directory that can be used to authenticate portal server users, IBM Dashboard Application Services Hub users, and optionally Netcool/OMNIbus Web GUI users. When a shared user registry is used, users are authenticated by the first server that they access and authentication tokens are passed to the other servers so that the user is not required to re-enter their credentials. A shared user registry is strongly recommended if you plan to use IBM Dashboard Application Services Hub with monitoring

dashboards. It allows you to take advantage of the authorization features supported by IBM Tivoli Monitoring and to enable single signon when the portal client is launched from IBM Dashboard Application Services Hub.

- Tivoli Performance Analyzer for predictive capability with Tivoli Monitoring so you can monitor resource consumption trends, anticipate future performance issues, and avoid or resolve problems more quickly.

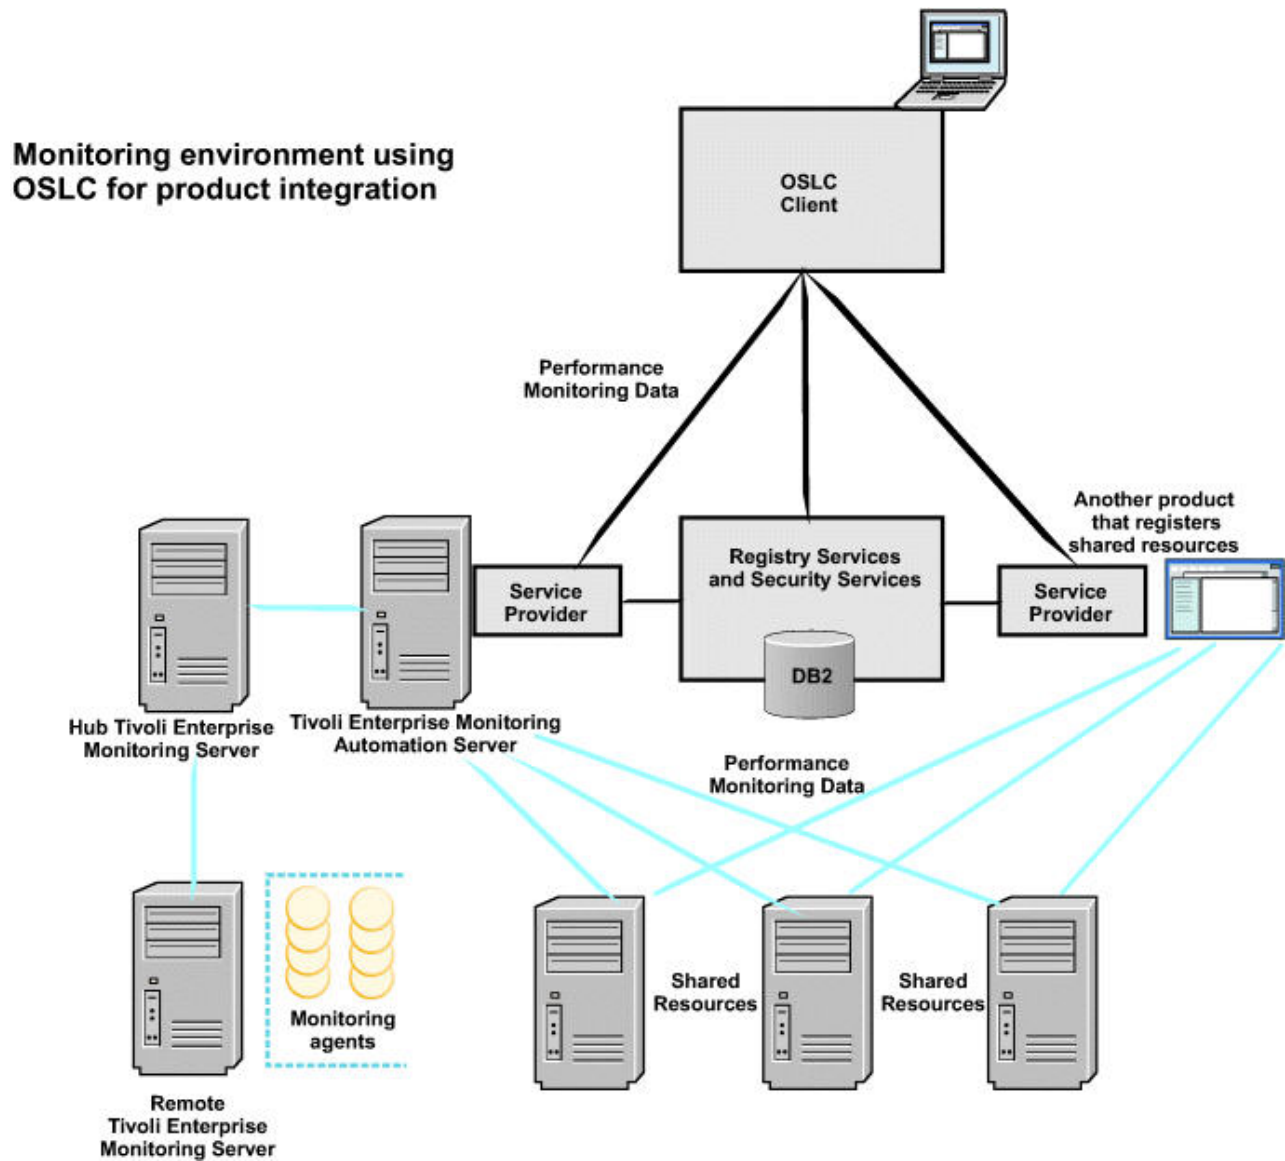


Figure 2. IBM Tivoli Monitoring environment using Open Services Lifecycle Collaboration for product integration

When IBM Tivoli Monitoring uses Open Services Lifecycle Collaboration and linked data principles for product integration, the IBM Tivoli Monitoring environment is extended by adding the following components:

- The Tivoli Enterprise Monitoring Automation Server is installed with the Hub monitoring server. It extends the Hub monitoring server by providing the Open Services Lifecycle Collaboration Performance Monitoring (OSLC-PM) service provider. The service provider registers monitoring resources such as computer systems, software servers, and databases with the Registry Services and also responds to HTTP GET requests for resource health metrics from OSLC clients.

- Other products such as Tivoli Application Dependency Discover Manager can also provide OSLC service providers that register shared resources such as computer systems, software servers, and databases with Registry Services and that respond to HTTP GET requests from OSLC clients.
- Registry Services is a Jazz for Service Management integration service that provides a shared data repository for products in an integrated service management environment. It reconciles resources registered by multiple service providers. OSLC client applications can retrieve a single record for a shared resource such as a computer system from Registry Services. The record contains URLs that the OSLC client application can use to retrieve additional details about the resource directly from each service provider using HTTP GET requests.
- Security Services is an optional Jazz for Service Management service that enables non-WebSphere based servers such as the Tivoli Enterprise Monitoring Automation Server to participate in Lightweight Third Party Authentication (LTPA) based single sign-on with OSLC clients installed on WebSphere® servers. For the V6.3 release of IBM Tivoli Monitoring, the Performance Monitoring service provider assumes that Registry Services and Security Services are installed into the same WebSphere Application Server.
- OSLC client applications are configured with the location of Registry Services. They query Registry Services to determine if a service provider has registered resources that they are interested in such as ComputerSystem resources. From Registry Services, the OSLC client also discovers the URLs that can be used to retrieve more information about a resource from the service providers that registered the resource. The OSLC client can then issue HTTP GET requests to the service provider to retrieve resource information in either RDF/XML, compact XML, or HTML format. Some OSLC clients can display the information retrieved from the service providers in a hover preview on the user interface so that the operator does not have to launch a separate application to see the details. For example, Tivoli Business Service Manager v6.1.1 can display hover preview in its service tree user interface to show health metrics for resources registered by the Performance Monitoring service provider and configuration and change history information registered for the same resources by Tivoli Application Dependency Discovery Manager version 7.2.1 FP4 or later.

Tivoli Enterprise Monitoring Server

The Tivoli Enterprise Monitoring Server is the key component on which all other architectural components directly depend. The monitoring server collects and controls all data and alerts received from monitoring agents.

The Tivoli Enterprise Monitoring Server (referred to as the *monitoring server*) is the first component to install to begin building the Tivoli Management Services foundation.

The monitoring server is the collection and control point for performance and availability data and alerts received from monitoring agents. It is also responsible for tracking the online or offline status of monitoring agents.

Because of the number of functions the monitoring server performs, large-scale environments usually include a number of monitoring servers to distribute the load. One of the monitoring servers is designated the hub monitoring server, and the remaining servers are termed remote monitoring servers. Each remote monitoring server must be located on its own computer and have a unique monitoring server name (node), but the architectures of various remote monitoring servers might differ from each other and from the hub monitoring server. In other words, a remote monitoring server running on UNIX can report to a hub monitoring server running on Windows.

The portal server communicates with the hub, which in turn controls the remote servers, as well as any monitoring agents that might be connected to it directly.

The monitoring server storage repository is a proprietary database format (referred to as the Enterprise Information Base, or EIB). The hub holds the master copy of the EIB, while the remote servers maintain a subset of the EIB relevant to them, which is synchronized with the hub.

Tivoli Enterprise Portal

Tivoli Enterprise Portal is the interface to your monitoring products. The Tivoli Enterprise Portal consists of the Tivoli Enterprise Portal Server and one or more clients.

The Tivoli Enterprise Portal Server (referred to as the *portal server*) manages data access through user workspace consoles (the portal clients). The portal server connects to a hub monitoring server; it retrieves data from the hub in response to user actions at a portal client, and sends the data back to the portal client for presentation. The portal server also provides presentation information to the portal client so that it can render the user interface views suitably.

The portal server also includes the optional dashboard data provider which is used by Dashboard Application Services Hub to retrieve read-only monitoring data from the hub monitoring server and monitoring agents. Dashboard Application Services Hub can display this data in dashboard applications such as the IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM or in custom dashboards. (The IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM applications are provided with the 7.2 release or later of IBM SmartCloud® Monitoring and IBM Tivoli Monitoring for Virtual Environments.)

The portal server uses a DB2, Derby, or Microsoft SQL database to store various artifacts related to presentation at the portal client.

The portal client provides access to the Tivoli Enterprise Portal. There are two kinds of portal client:

- Browser client interface (automatically installed with Tivoli Enterprise Portal Server): You can run the browser client using Microsoft Internet Explorer or Mozilla Firefox; it connects to a Web server running in the Tivoli Enterprise Portal Server. Running the browser client is supported only on Windows and Linux computers.
- Desktop client interface: A Java-based graphical user interface on a Windows or Linux workstation. After the desktop client is installed and configured, you can use it to start Tivoli Enterprise Portal in desktop mode. You can also download and run the desktop client using Java Web Start.

“Java Web Start clients” on page 347

No post-installation configuration of the JRE is required if the Java Web Start is used to deploy the Tivoli Enterprise Portal client.

“Configuring clients, browsers, and JREs” on page 342

The configuration required for Tivoli Enterprise Portal clients depends upon the client deployment mode being used, the browser being used, the Java runtime environment (JRE) being used, and the operating system the client is being used on.

Tivoli Enterprise Monitoring Agents

Monitoring agents are data collectors. Agents monitor systems, subsystems, or applications, collect data, and pass the data to the Tivoli Enterprise Portal through the monitoring server.

An agent interacts with a single system or application and, in most cases, is located on the same computer where the system or application is running.

There are two types of monitoring agents:

- Operating system (OS) agents that monitor the availability and performance of the computers in your monitoring environment. An example of an OS agent is the Monitoring Agent for Windows OS, which monitors Windows XP and Windows 2003 operating systems.

A special type of operating system agent, the agentless monitor, is also available. It enables a remote node to monitor the health of nonessential desktop operating systems via a standard monitoring API such as SNMP and thus is also called a remote OS agent.

In addition, there is another class of operating-system agents, the System Monitor Agent. These lighter-weight agents (they require a much smaller footprint than full-function Tivoli Monitoring OS agents) are configured locally to the agent node. This configuration enables them to be deployed autonomously (in other words, without the support of a Tivoli Enterprise Monitoring Server): they send EIF or SNMP event information directly to an event server such as Netcool/OMNIBus. The System Monitor Agents are meant as a replacement for the OMNIBus System Service Monitor agents.

- Other agents (referred to as *application agents* or *non-OS agents*) that monitor the availability and performance of systems, subsystems, and applications. An example of a non-OS agent is IBM Tivoli Monitoring for Microsoft Exchange, which monitors the Microsoft Exchange Server.

The Performance Analyzer Warehouse Agent adds predictive capability to Tivoli Monitoring so you can monitor resource consumption trends, anticipate future performance issues, and avoid or resolve problems more quickly. For more information about the Performance Analyzer Warehouse Agent, see the *IBM Tivoli Monitoring Performance Analyzer User's Guide*.

You can also create your own IBM Tivoli Monitoring agent via the Agent Builder, a set of tools for creating agents and adding value to existing agents. Using the Agent Builder, you can quickly create, modify, and test an agent to collect and analyze data about the state and performance of different resources, such as disks, memory, CPU, or applications. The builder creates a data provider that allows you to monitor three types of data:

Availability

Process and service availability and functionality tests

Windows event log

Specific information from the Windows Event Log

External data sources

Data from external sources such as Windows Management Instrumentation (WMI), Performance Monitor (PerfMon), Simple Network Management Protocol Version 1 (SNMP V1), external scripts, and log files

With the Agent Builder's customizable Graphical User Interface installer, you can create agent-installation packages for easy agent distribution. A key feature of the installer is its ability to package and distribute extensions to existing agents. This lets you develop new situations, queries, and workspaces for an existing IBM Tivoli Monitoring V6.x agent. For complete information about the Agent Builder, see the *Agent Builder User's Guide*.

In most cases the recommended choice for customized agents is the Agent Builder.

The *IBM Tivoli Monitoring Base DVD* (in addition to providing the Tivoli Enterprise Monitoring Server and its application support, Tivoli Enterprise Monitoring Automation Server, the Tivoli Enterprise Portal Server and its application support, and the Tivoli Enterprise Portal desktop and browser clients together with their application support) also contains the Tivoli Data Warehouse agents: the Warehouse Proxy Agent and the Summarization and Pruning Agent, and the Tivoli Performance Analyzer. This DVD is platform-specific (Windows, Linux, or UNIX).

Use the *IBM Tivoli Monitoring Agents DVD* to install the monitoring agents in the following list (as well as agentless OS monitoring agents). Note that this DVD, however, is platform-nonspecific; that is, it applies to Windows, Linux, and UNIX environments.

IBM i
Windows OS
Linux OS
UNIX OS

Note: For z/OS customers, IBM also provides a family of OMEGAMON Monitoring Agents that monitor both the z/OS operating system (as well as its key subsystems: VTAM®, CICS®, IMS™, DB2, and storage subsystems) and the z/VM® operating system (as well as any Linux guests running under it). A complete suite of OMEGAMON product documentation is provided in the IBM Tivoli Monitoring Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/welcome.htm) information center.

See “Selecting the correct support media” on page 277 for information on what media to use for installing application support for distributed agents.

Tivoli Data Warehouse

With Tivoli Data Warehouse, you can analyze historical trends from monitoring agents. The Tivoli Data Warehouse uses a DB2 Database for Linux, UNIX, and Windows, DB2 on z/OS, Oracle, or Microsoft SQL database to store historical data collected across your environment.

You can generate warehouse reports for short-term or long-term data through the Tivoli Enterprise Portal. Warehouse reports provide information about the availability and performance of your monitoring environment over a period of time. Monitoring agents such as the OS agents also provide predefined reports for Tivoli Common Reporting, a Jazz for Service Management component. In addition, Tivoli Common Reporting has a custom reports capability. You can also use third-party warehouse reporting software, such as Crystal Reports or Brio, to generate reports.

Three specialized agents interact with the Tivoli Data Warehouse:

- The Warehouse Proxy Agent receives data collected by monitoring agents and moves it to the Tivoli Data Warehouse database.
- The Summarization and Pruning Agent provides the ability to customize the length of time for which to save data (pruning) and how often to aggregate granular data (summarization) in the Tivoli Data Warehouse database.
- The Tivoli Performance Analyzer extends the capability of Tivoli Monitoring by analyzing and enriching the data that is collected by its monitoring agents and by providing reports about the performance and capacity of your systems.

The Warehouse Proxy Agent, the Summarization and Pruning Agent, and the Tivoli Performance Analyzer, and support for those agents, are installed from the *IBM Tivoli Monitoring Base DVD*.

Related concepts:

Chapter 19, “Tivoli Data Warehouse solutions,” on page 509

The term *Tivoli Data Warehouse solution* refers to the set of IBM Tivoli Monitoring components, successfully installed and configured, that interact to collect and manage historical data.

Event synchronization component

The event synchronization component, the Event Integration Facility or EIF, sends updates to situation events that have been forwarded to a Netcool/OMNIBus ObjectServer or a Tivoli Enterprise Console event server back to the monitoring server.

Figure 3 on page 14 shows the flow of situation events from the hub monitoring server to the event server as EIF events and the flow of updates to the situation events back to the hub monitoring server. The Situation Event Console and the Common Event Console in the Tivoli Enterprise Portal client, event views in monitoring dashboards on IBM Dashboard Application Services Hub, and the event server's event views are synchronized with the updated status of the events.

If you are monitoring event data from a supported event management system in the Tivoli Enterprise Console event view or the Common Event Console view, you can filter out forwarded events. For more information, see the *IBM Tivoli Monitoring Administrator's Guide*.

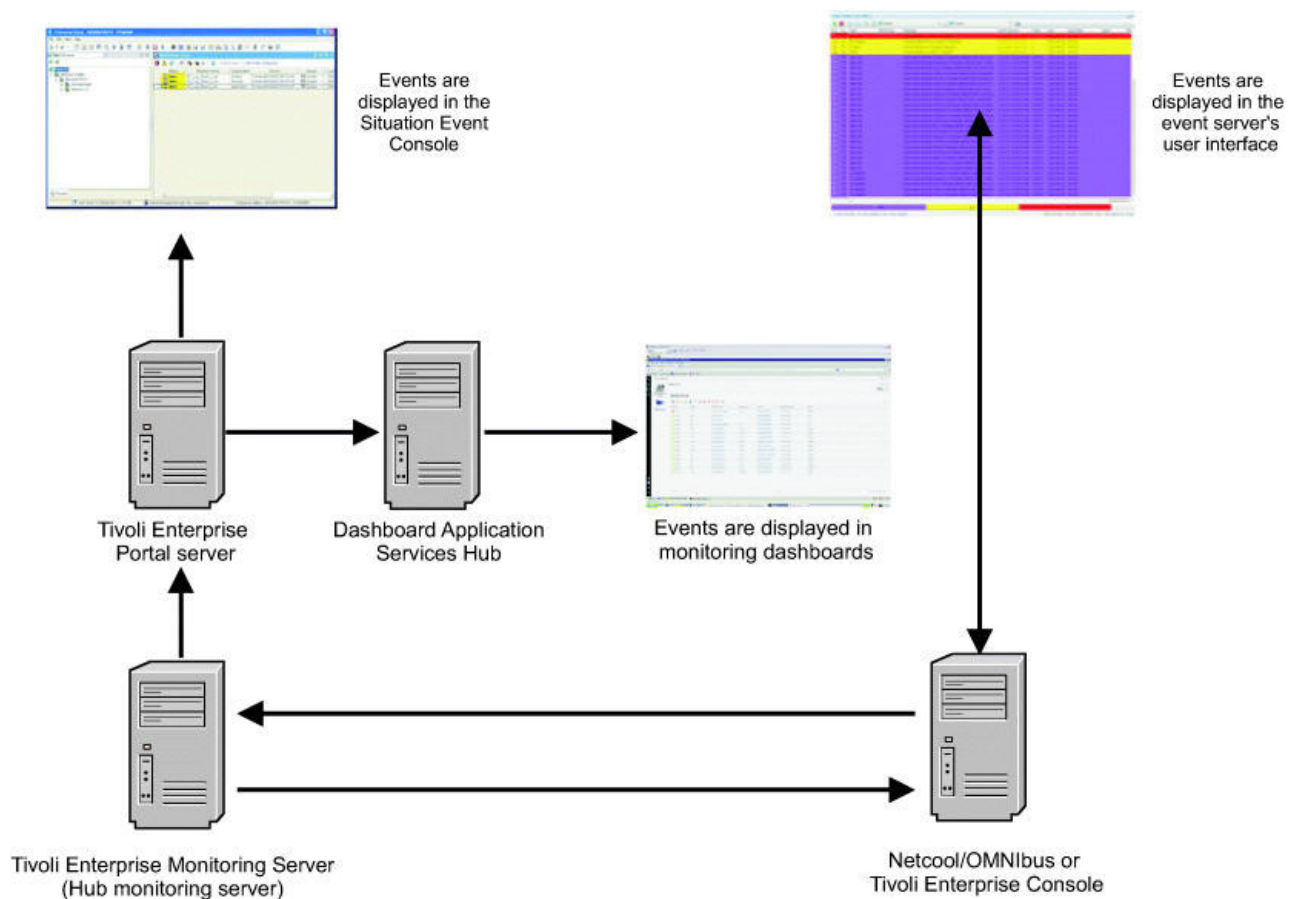


Figure 3. Event synchronization overview

Related concepts:

Chapter 29, “Integrating event management systems,” on page 695

If you are using either Netcool/OMNIBus or Tivoli Enterprise Console in addition to IBM Tivoli Monitoring to manage events in your enterprise, you can integrate and manage events from a single console. You can integrate event management by forwarding events reported by Tivoli Enterprise Monitoring Agents to either event system for correlation and management - changes in event status made on the event system are reflected back to the hub monitoring server that forwarded them. Or you can enable events reported by a Tivoli Enterprise Monitoring Agent to be passed directly to OMNIBus for processing, thereby bypassing the monitoring server entirely.

Tivoli Enterprise Portal Server extended services

Tivoli Enterprise Portal Server extended services (TEPS/e) is an embedded, co-located extension of the Tivoli Enterprise Portal Server that provides Java EE-based Application Server integration facilities.

TEPS/e supports a federated user repository such as those based on the Lightweight Directory Access Protocol (LDAP). For more information, see the *IBM Tivoli Monitoring Administrator's Guide*.

tacmd Command Line Interface

You can use the command line interface (CLI) in IBM Tivoli Monitoring to perform administrative tasks such as creating and managing a situation or starting and stopping an agent.

Tivoli administrative commands (tacmd) serve both UNIX-based and Windows operating systems. You do not need to launch or set up the shell environment. For example, you can enter the command `tacmd help` from any directory on the computer that hosts the monitoring server. The tacmd commands provide the following capabilities, among others:

- Obtain details about installed IBM Tivoli Monitoring components.
- Stop, start, restart, and configure monitoring agents.
- Create a node for deployment of monitoring agents.

For the complete list of tacmd commands, see the *IBM Tivoli Monitoring Command Reference*.

IBM Dashboard Application Services Hub and dashboard applications

The Dashboard Application Services Hub is a web-based console component that provides common task navigation, aggregation of data from multiple products into a single view, and message passing between views from different products. The Dashboard Application Services Hub has a core set of components that provide such administrative essentials as network security.

The IBM Tivoli Monitoring dashboard data provider retrieves monitoring agent data for display in the Dashboard Application Services Hub. The dashboard data provider is optionally enabled during the Tivoli Enterprise Portal Server configuration. With the dashboard data provider enabled, Dashboard Application Services Hub users can retrieve read-only data from the hub monitoring server and monitoring agent for display in monitoring dashboards such as the IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM or in custom dashboards. (The IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM applications are provided with the 7.2 release or later of IBM SmartCloud Monitoring and IBM Tivoli Monitoring for Virtual Environments.)

IBM Infrastructure Management Dashboards for Servers is provided with IBM Tivoli Monitoring and allows Dashboard Application Services Hub users to display events for monitoring agents, managed system groups, and key Linux OS agent, UNIX OS agent, and Windows OS agent health metrics.

Because monitoring dashboard users must be authenticated by Dashboard Application Services Hub and by the Tivoli Enterprise Portal Server, you should configure each of these servers to use the same LDAP user repository for authentication and also enable single sign-on.

Tivoli Authorization Policy Server

The Authorization Policy Server allows you to define roles and permissions, which control the access that dashboard users have to managed systems and managed system groups displayed in monitoring dashboards of the IBM Dashboard Application Services Hub.

The Authorization Policy support consists of two installable components:

- Tivoli Authorization Policy Server
- tivcmd Command Line Interface for Authorization Policy

The Authorization Policy Server must be installed on the same system as an IBM Dashboard Application Services Hub. The tivcmd Command Line Interface is used by administrators to create and work with roles and permissions. It can be installed on the same system as the Authorization Policy Server, or it can be installed on an administrator's computer and communicate via HTTP or HTTPS to the Authorization Policy Server.

The Authorization Policy Server and the tivcmd CLI follow the role-based access control (RBAC) standard. According to this standard, roles are created for job functions, and permissions are assigned to roles. RBAC simplifies common security administration functions such as adding users or changing a

user's department because individual users are not assigned permissions directly. Instead, a user acquires permissions based on the role (or roles) that user belongs to.

The Authorization Policy Server provides more granular access control than Tivoli Enterprise Portal user management permissions and application assignments. For example, Tivoli Enterprise Portal user management allows you to assign application types such as Windows OS to users or user groups. This assignment allows those users or user groups to view all Windows OS agent managed systems in the Tivoli Enterprise Portal client or in the Dashboard Application Services Hub. However, the Authorization Policy Server can be used to create roles and permissions that allow users or user groups to view data from either specific managed systems or from any managed system in specific managed system groups. For example, you can use Tivoli Enterprise Portal or `tacmd` CLI to create a managed system group called `easternregion` that only contains Windows OS agents at your eastern region data center. Then you can use the `tivcmd` CLI to create a role that has permission to view the managed systems in the `easternregion` managed system group. When users or user groups are assigned to this role then they will only see the `easternregion` managed system group and its members in the Dashboard Application Services Hub.

Enforcement of authorization policies occurs in the dashboard data provider component of the Tivoli Enterprise Portal Server whenever a dashboard user attempts to view a situation event, or to display attribute data from a monitoring agent or managed system group. To avoid delays caused by querying the master policy file store across the network, the dashboard data provider downloads its own local copy of the policy file store from the Authorization Policy Server and uses it for authorization checking.

When a V6.3 Tivoli Enterprise Portal Server is installed, authorization policy enforcement is disabled by default. When it is disabled, the dashboard data provider uses Tivoli Enterprise Portal user management permissions and application assignments to determine what agent types the dashboard users are allowed to view. At a later time, after your dashboard environment is working and your administrators have defined an appropriate set of roles and permissions, you can enable authorization policy enforcement by reconfiguring the Tivoli Enterprise Portal Server. Once authorization policy enforcement is enabled in the portal server, the dashboard data provider uses the authorization policies instead of Tivoli Enterprise Portal permissions and application assignments to determine what monitoring resources a dashboard user can access.

Authorization policies only control what monitoring resources a user can access in Dashboard Application Services Hub. If a dashboard user also uses the Tivoli Enterprise Portal, they might see a different set of resources because the portal client uses Tivoli Enterprise Portal user management permissions and application assignments to determine what resources a user can access.

tivcmd Command Line Interface for Authorization Policy

The `tivcmd` Command Line Interface for Authorization Policy has a similar look and feel to the existing `tacmd` Command Line Interface. An ITM administrator uses the `tivcmd` Command Line Interface to log in to the Authorization Policy Server and then uses the set of `tivcmd` commands such as **`createrole`**, **`grant`**, **`exclude`**, and **`addtorole`** to construct authorization policies that consist of roles and permissions that control which job functions have access to which IBM Tivoli Monitoring resources.

The authorization policies are saved in a master policy file store that is maintained on the Authorization Policy Server system. The `tivcmd` Command Line Interface can be located on the same system as the Authorization Policy Server, or it can be installed on a separate system and communicate via HTTP or HTTPS to the Authorization Policy Server.

For the full syntax of `tivcmd` CLI commands and parameter descriptions, see the *IBM Tivoli Monitoring Command Reference*.

Tivoli Enterprise Monitoring Automation Server

The Tivoli Enterprise Monitoring Automation Server is an optional component that can be installed on the same system as the hub monitoring server. It extends the functionality of the hub monitoring server.

The automation server includes the Open Services Lifecycle Collaboration Performance Monitoring (OSLC-PM) service provider. The Performance Monitoring service provider registers resource types such as Computer System, Software Server, Database, Software Module and Service Instance on behalf of monitoring agents. These resource types are defined using the Common Resource Type Vocabulary (CRTV). Agents provide a template that maps their monitoring data to CRTV resources. The template is installed with the monitoring server application support of the agent. Products such as IBM Tivoli Monitoring that discover and manage shared IT resources can register these IT resources and the services they offer with Registry Services. Other products can consume data by querying Registry Services for the managed resources or the associated service providers of interest. In the query responses, Registry Services returns links to the service providers that can be queried to retrieve more information about the managed resources.

The Performance Monitoring service provider also supports the OSLC-PM RESTful API for retrieving linked data about monitored resources. It accommodates HTTP GET requests for the RDF/XML, compact XML and HTML content types. When RDF/XML and HTML content is requested, the API returns resource metrics defined by the OSLC-PM domain and the IBM Tivoli Monitoring private namespace. For more information about OSLC and Loosely Coupled Integration, see the following links:

- OSLC community
- Performance Monitoring working group
- IBM Tivoli Monitoring OSLC private namespace schema
- Reconciliation working group and Common Resource Type Vocabulary
- Loosely coupled integration at ISM Connect

Security Services is an optional Jazz for Service Management component that enables non-WebSphere based applications such as the Performance Monitoring service provider to participate in LTPA based single sign-on. It should be installed and configured if you want the Performance Monitoring service provider to authenticate requests it receives from OSLC clients. This is the only authentication method supported by the service provider. The application servers for OSLC clients also must be configured for single sign-on with Security Services and include an LTPA token in HTTP GET requests sent to the service provider. When the Performance Monitoring service provider receives an HTTP GET request from an OSLC client, it forwards the LTPA token to Security Services to authenticate the request. If the request does not contain a LTPA token or Security Services indicates that the token is not valid or has expired, the Performance Monitoring service provider returns an HTTP status code indicating that the request could not be authenticated. For the V6.3 release of IBM Tivoli Monitoring, the Performance Monitoring service provider assumes that Registry Services and Security Services are installed into the same WebSphere Application Server.

For more information, see the *Performance Monitoring service provider* topic in the *IBM Tivoli Monitoring Administrator's Guide* and "Installing and configuring Tivoli Enterprise Monitoring Automation Server" on page 219.

For more information on Registry Services and Security Services, see the Jazz for Service Management Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/psc_ic-homepage.html).

Tivoli Performance Analyzer

Tivoli Performance Analyzer adds predictive capability to Tivoli Monitoring so you can monitor resource consumption trends, anticipate future performance issues, and avoid or resolve problems more quickly.

For example, you can use Tivoli Performance Analyzer to predict application bottlenecks and create alerts for potential service threats.

Tivoli Performance Analyzer helps IT managers to answer the following questions so they can optimize IT capacity:

- When will my application fail to meet service levels?
- How will application performance change if I modify the infrastructure?
- What is the best hardware solution to meet my performance and cost goals?
- Where are my under-utilized servers and networks?
- Which servers and network components really need an upgrade?
- Which application will experience the next performance issue? When?

Providing accurate IT forecasts and appropriate IT capacity to meet business needs are two major goals in the capacity management area. You can use the following key performance indicators (KPIs) to measure the critical success factors:

- Total value of unplanned or unused capacity expenditures
- Percent of capacity forecasts that were accurate
- Number of inaccurate business forecast inputs provided
- Number of incidents related to capacity or performance issues

Capacity management

Capacity management is a forward looking process that aims to align system capacity to the current and future demand in a cost-efficient manner. Within capacity management, performance analysis is a key activity. It requires tools and methods that predict how IT resources will behave in the near and mid term so you can avoid incidents and problems rather than have to react to them.

Tivoli Performance Analyzer extends the capability of Tivoli Monitoring by analyzing and enriching the data that is collected by its monitoring agents and by providing reports about the performance and capacity of your systems. Tivoli Performance Analyzer performs the following functions for individual IT components:

- Gathers and stores IT components such as CPU, disk, and memory utilization in a central data repository
- Provides a predictive analysis component that indicates trends in IT component utilization
- Retains the analyzed data in a central repository for reporting purposes and for input to the component management process

Chapter 2. Pre-deployment phase

This chapter contains information that assists you in assessing your environment and planning the deployment of product components.

The successful use and availability of your IBM Tivoli Monitoring environment is the result of a well-planned and executed deployment. Tivoli Monitoring products use a set of service components (known collectively as Tivoli Management Services) that are shared by a number of other product suites, including IBM Tivoli OMEGAMON XE monitoring products, IBM Tivoli Composite Application Manager products, System Automation for z/OS, Web Access for Information Management, and others. Much of the information in this guide is also relevant to these products.

Planning checklist

Correctly planning your installation is probably the most important thing that you can do to achieve a smoothly running environment. By carefully planning the size of your servers, understanding your network topology, and understanding your key requirements, you can create a monitoring environment that meets your needs.

Careful planning is a team effort. Engage the networking team to ensure you understand all of the network firewall configurations and to open ports in the firewall. Understand the key requirements such as fault tolerance, executive dashboards, and reporting. The following sections outline the key planning items that you must complete to have a successful deployment.

Use the Planning checklist in Table 1 to ensure that all important planning tasks are accomplished. Perform all of these tasks before starting your Tivoli Monitoring installation. The following sections of this guide provide the necessary information to complete this checklist.

Table 1. Planning checklist

Planning activity	Comments	Status
Attend IBM Tivoli Monitoring Administrator Training.	Sample: 2 people attended Administrator class	Sample: Complete
Determine the desired platforms for Tivoli Monitoring infrastructure components.	See Chapter 3, "Hardware and software requirements," on page 75.	
Identify the number of Tivoli Monitoring agents (typically 3 monitoring agents per monitoring server).		
Determine high availability and disaster recovery requirements. Map those requirements to your Tivoli Monitoring deployment.		
Determine whether a firewall gateway is required and between which components.		
Determine the location of any remote monitoring servers.		
Determine that an adequate number of Warehouse Proxy Agents are planned.		
Provide Network administrators with a list of ports to open.		

Table 1. Planning checklist (continued)

Planning activity	Comments	Status
Complete the Warehouse load projection spreadsheet. See “Locating and sizing the Warehouse Proxy Agent” on page 37.		
Determine the hardware required for your Tivoli Monitoring components.		
Download all the required software for your supported operating systems including all monitoring agents.		
Complete your deployment Verification Test Plan.		
Review the <i>Readme and Documentation Addendum</i> and determine the installation steps based on your platform.		
Determine the method of user authentication. See paragraphs below.		
If you plan to use Netcool/OMNIBus as your event server, work with your Netcool/OMNIBus administrator to determine the location of the Netcool/OMNIBus components and to plan the installation of the IBM Tivoli Monitoring Event Synchronization component.		
<p>If you plan to use the IBM Infrastructure Management Dashboards for Servers application or other monitoring dashboard applications for the IBM Dashboard Application Services Hub:</p> <ul style="list-style-type: none"> • determine where to install the IBM Dashboard Application Services Hub and its prerequisites. • determine if you will setup a network repository that can be used for installing Jazz for Service Management and IBM Tivoli Monitoring components throughout your enterprise or if you will use the installation launchpad applications provided by Jazz for Service Management and IBM Tivoli Monitoring to perform installations directly from the installation media. • determine whether the Authorization Policy Server or Tivoli Enterprise Portal user management will be used to authorize which resources can be viewed by dashboard operators. • if you plan to use the Authorization Policy Server, determine where the tivcmd Command Line Interface for Authorization Policy component should be installed. Administrators who create and work with authorization policies will use this CLI. 		
If you plan to use IBM Tivoli Monitoring reports with Tivoli Common Reporting, determine where to install Tivoli Common Reporting and its prerequisites.		

Table 1. Planning checklist (continued)

Planning activity	Comments	Status
If you plan to use the OSLC Performance Monitoring service provider, determine where to install the Jazz for Services Management Registry Services component and, optionally, the Security Services component and their prerequisites.		

Tivoli Monitoring V6.3 supports integration with LDAP user registries for authentication. See “Security options” on page 137 for more information.

LDAP SSL requires some actions by an LDAP administrator that are not covered by the Tivoli Monitoring V6.3 documentation. Here are some LDAP SSL Web pages for working with LDAP servers:

- Configuring Microsoft Active Directory for SSL access
- Configuring Oracle Java System Directory Server for SSL access
- Configuring the Tivoli Directory Server client for SSL access

If the hub monitoring server is configured to use an LDAP server, LDAP SSL communication will also require the creation of a GSKit keystore for the hub monitoring server; see the *Securing communication* chapter in the *IBM Tivoli Monitoring Administrator's Guide*.

Understanding Tivoli Monitoring and your network

Consider your network infrastructure when you are planning your Tivoli Monitoring deployment.

For all successful deployments, you must understand the network infrastructure. When planning your Tivoli Monitoring deployment, take into account each of the following important factors, which are discussed in further detail in this guide:

- Locations of firewalls
- Whether you have NAT (network address translation)
- Network bandwidth between WAN (wide area network) links
- Number of agents that are located across WAN links

Tivoli Monitoring has several options for the communication protocols that are used by the product components that communicate with a monitoring server. The IP PIPE and IP SPIPE are the protocols used when firewalls are used in the environment. Unlike the IP (UDP) protocol, the communications for IP PIPE and IP SPIPE can be restricted to a single port.

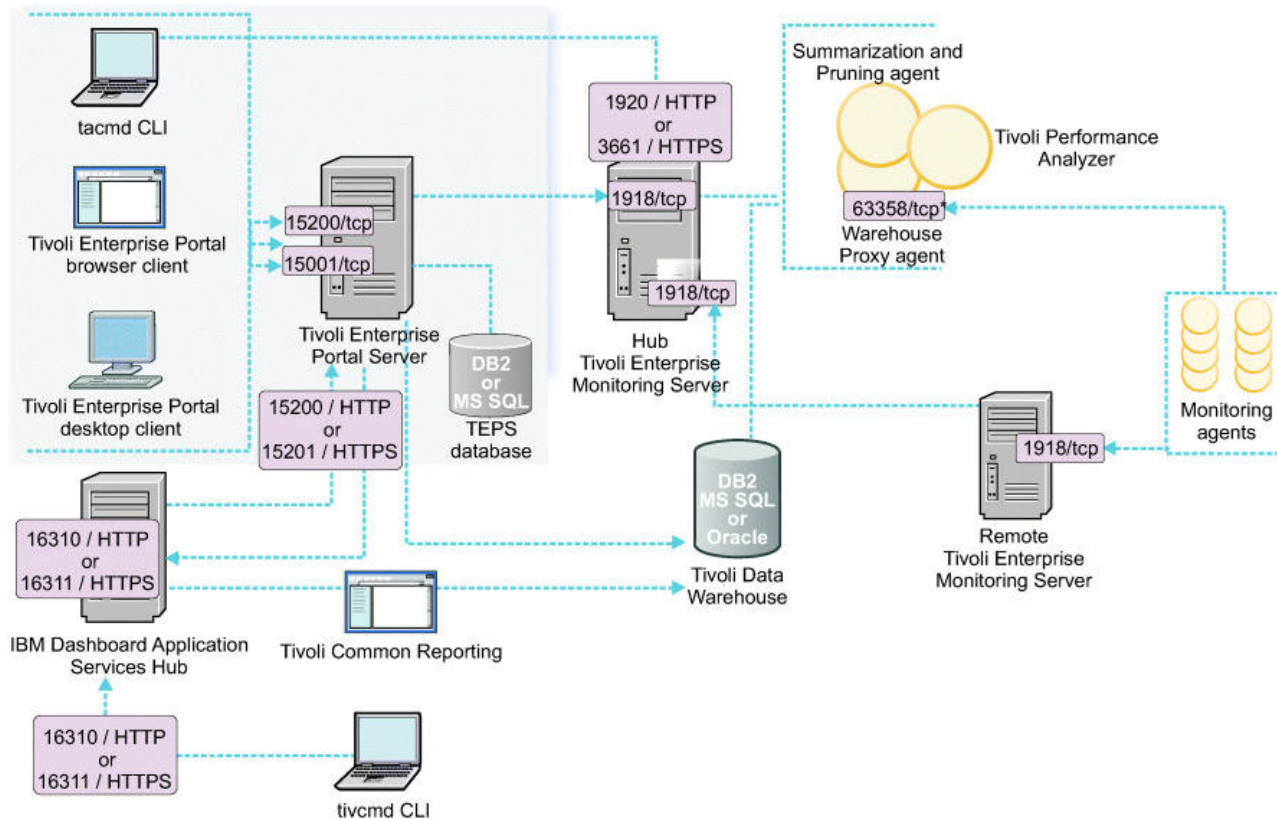


Figure 4. Tivoli Monitoring V6.3 communications model

Note:

1. The asterisk '*' in 63358/tcp* under Warehouse Proxy Agent means the port is configurable (by using the SKIP and COUNT parameters for the Warehouse Proxy Agent).
2. The default port for SSL communication with a monitoring server is 3660/TCP. (For non-SSL communication the default port is 1918/TCP.)
3. Figure shows the default port numbers for the Dashboard Application Services Hub but it is possible to configure the dashboard hub to use different port numbers.

Determine if you require a firewall gateway

For most environments, using the firewall gateway is not required when deploying the Tivoli Monitoring software. However, in some cases, the firewall gateway is the only way to traverse the complex firewalls in a network.

The following section describes the scenarios when the firewall gateway is required. In addition, the section outlines the optimal locations for the firewall gateway.

Use the firewall gateway for *any* of the following scenarios:

- A single TCP connection cannot be made to span between Tivoli Monitoring components. One example is when there are multiple firewalls between these components and a policy that does not allow a single connection to traverse multiple firewalls.
- Connection requirements do not allow the Tivoli Monitoring default pattern of connections to the hub Tivoli Enterprise Monitoring Server. One example is when agents that are located in a less-secure zone connect to the monitoring server located in a more-secure zone. Security policy allows a connection to be established only from a more-secure zone to a less-secure zone, but not the other way round.

- You must reduce open firewall ports to a single port or connection. For example, rather than opening the port for every system being monitored, consolidate the ports into a single concentrator.
- You must manage agent failover and Tivoli Enterprise Monitoring Server assignment symbolically at the hub monitoring server end of the gateway. Because gateway connections are made between matching service names, an administrator can change the failover and monitoring server assignment of downstream gateway agents by changing the client proxy bindings next to the hub monitoring server.

Network address translation (NAT) alone is not a reason to use the firewall gateway, which is content-neutral and can proxy any TCP connection. In most cases, NAT processing can be handled by the PIPE protocol (IP.PIPE or IP.SPIPE) without the firewall gateway.

For detailed information on installing and configuring the firewall gateway, see Appendix C, “Firewalls,” on page 861.

Determine where to place your Tivoli Monitoring components

There are a few factors to consider when determining the location of your Tivoli Monitoring components. The primary factors are firewalls and slow network connections. Before discussing the locations of these components in the data center, you must understand these components, the roles that they play, and what affects the load on these components.

The Tivoli Monitoring components and data paths to the Tivoli event management products are illustrated in Figure 5.

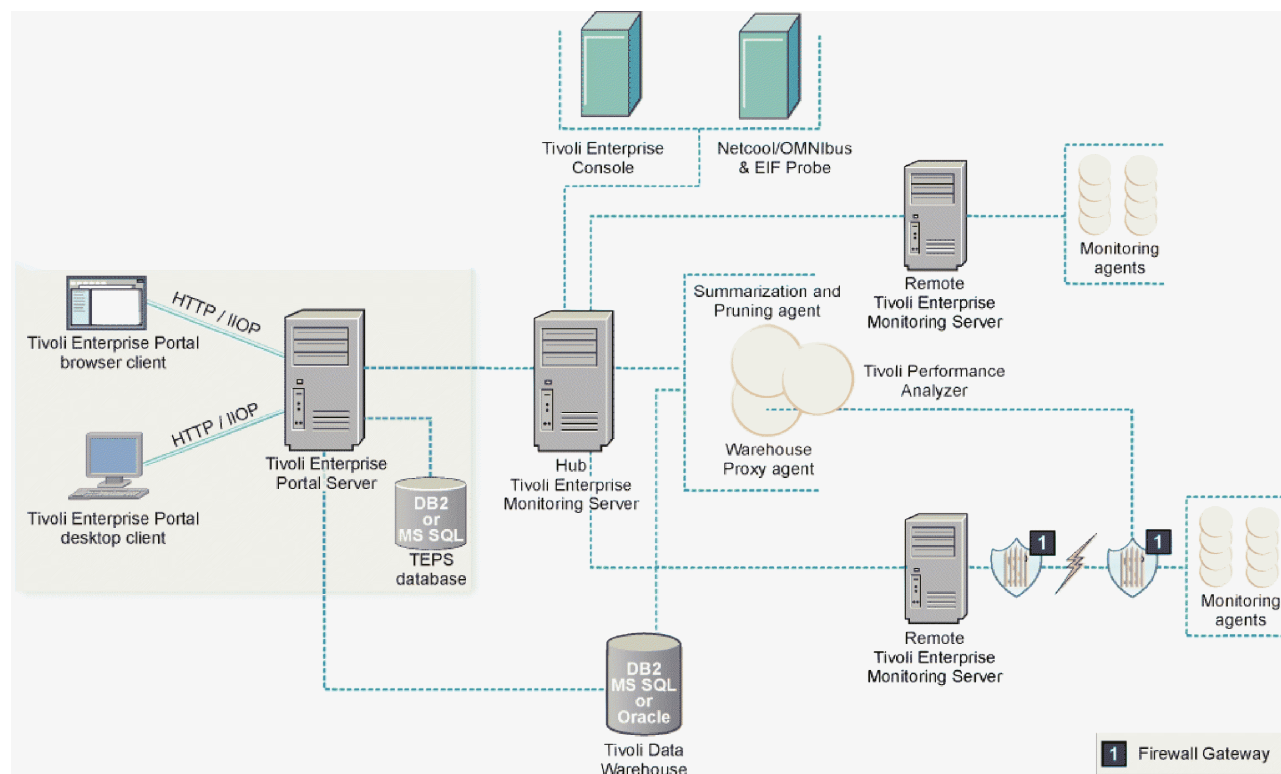


Figure 5. Tivoli Monitoring component architecture including firewall gateway

Tivoli Monitoring installation requires the following optional and mandatory components:

- “Tivoli Enterprise Monitoring Server” on page 24
- “Tivoli Enterprise Portal Server” on page 25
- “Tivoli Enterprise Portal client” on page 27

- “tacmd Command Line Interface” on page 28
- “Warehouse Proxy Agent” on page 29
- “Warehouse Summarization and Pruning Agent” on page 30
- “Tivoli Data Warehouse” on page 30
- “Dashboard Application Services Hub with Tivoli Common Reporting” on page 30
- “Netcool/OMNibus integration” on page 31
- “Tivoli Enterprise Console integration” on page 31
- “Dashboard Application Services Hub with dashboard applications and Authorization Policy Server” on page 31
- “tivcmd Command Line Interface for Authorization Policy” on page 32
- “Tivoli Enterprise Monitoring Automation Server, Performance Monitoring service provider, Registry Services, and Security Services” on page 32
- “IBM Tivoli Agent Builder” on page 33
- “Firewall gateway” on page 33
- “Tivoli Performance Analyzer” on page 33
- “Monitoring agent for IBM Tivoli Monitoring 5.x Endpoint” on page 34

The specific hardware and software prerequisites for each of these components are listed in Chapter 3, “Hardware and software requirements,” on page 75.

Tivoli Enterprise Monitoring Server

The Tivoli Enterprise Monitoring Server, referred to as the monitoring server, is the first component installed to begin building the IBM Tivoli Monitoring Services foundation. The monitoring server is the key component on which all other architectural components directly depend.

The monitoring server is the collection and control point for alerts received from agents. The monitoring server collects the performance and availability data of the agents.

The monitoring server is also responsible for processing *heartbeats* to track the online or offline status of monitoring agents. A monitoring agent sends a heartbeat every 10 minutes (this is configurable). If a heartbeat signal is not received from an agent within the heartbeat interval (plus a 3-minute grace period), the monitoring server considers the monitoring agent to be offline. The online or offline status of the monitoring agent is also called the *managed system status*.

In addition to the roles described in the previous paragraphs, the hub monitoring server also provides the following functions:

- Distributes situations down to the monitoring agents.
- Used for remote deployment and remote configuration and control of monitoring agents.
- Responsible for the processing of CLI and SOAP requests that can be used by users for automation of their Tivoli Monitoring environment.

There are two types of monitoring servers: the hub monitoring server and the remote monitoring server. The hub monitoring server is the focal point for the entire Tivoli Monitoring environment. The hub monitoring server is under a significant load. Work on the hub includes connections from the remote monitoring server, authentication, situations, policies, and workflows.

The hub monitoring server stores, initiates, and tracks all situations and policies and is the central repository for storing all active conditions and short-term data on every Tivoli Enterprise Monitoring Agent (monitoring agent). The hub monitoring server is also responsible for initiating and tracking all generated Take Action commands.

The monitoring server storage repository is a proprietary database format, referred to as the Enterprise Information Base (EIB), that is grouped as a collection of files located on the monitoring server. These files start with a file name prefix `qa1` and are located in the following directories:

- **Windows** `installation_dir\cms`
- **Linux** **UNIX** `installation_dir/tables/tems_name`

where `installation_dir` specifies the Tivoli Monitoring installation directory and `tems_name` specifies the Tivoli Enterprise Monitoring Server name.

Note: You can use the `CANDLEHOME` command on UNIX or Linux and `CANDLE_HOME` on Windows system to locate the home directory for Tivoli Monitoring.

Place the hub monitoring server inside the data center on a high-performance network (100 Megabits per second or higher). Connectivity between the Tivoli Enterprise Portal Server and hub monitoring server as well as between the hub monitoring server and most of the remote monitoring server must be fast and reliable. For large environments, use a multi-processor server. Some hardware configuration guidance is identified in “Sizing your Tivoli Monitoring hardware” on page 34.

The remote monitoring server is a collection point for the agents that are connected to that remote monitoring server. Certain types of situations run on the remote monitoring server. The load on the remote monitoring server is typically low. Load is driven higher if historical data is collected at the remote monitoring server instead of at the agents.

Placement of the remote monitoring server depends on a few factors. Plan your firewalls early to ensure that communications can be established between the Tivoli Monitoring components with only a modest number of holes in the firewall.

By locating the Warehouse Proxy Agent on the same computer as the remote monitoring server, you can negotiate NAT environments with their historical data collection. For remote locations connected over a slow network, place the remote monitoring server at the remote location if there are significant numbers of computers. For remote locations with just a few computers, it doesn't make sense to place a remote monitoring server at the remote location.

Tivoli Enterprise Portal Server

The Tivoli Enterprise Portal Server, referred to as the portal server, is a repository for all graphical presentation of monitoring data. The portal server database consists of all user IDs and user access controls for the Tivoli Enterprise Portal monitoring workspaces and optionally for monitoring dashboard applications installed into Dashboard Application Services Hub. The portal server provides the core presentation layer, which allows for retrieval, manipulation, analysis, and pre-formatting of data. The portal server manages data access through user workspace consoles for the Tivoli Enterprise Portal.

The portal server also includes the optional dashboard data provider which is used by Dashboard Application Services Hub to retrieve read-only monitoring data from the hub monitoring server and monitoring agents. Dashboard Application Services Hub can display this data in dashboard applications such as the IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM, or in custom dashboards. (The IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM applications are provided with the 7.2 or later release of IBM SmartCloud Monitoring and IBM Tivoli Monitoring for Virtual Environments.) The dashboard data provider uses either authorization policies created using the `tivcmd` Command Line Interface and the Authorization Policy Server or Tivoli Enterprise Portal permissions to control view access to monitoring resources in dashboard applications.

The portal server keeps a persistent connection to the hub monitoring server, and can be considered a logical gateway between the hub monitoring server and the Tivoli Enterprise Portal client (portal client)

and IBM Dashboard Application Services Hub. Any disconnection between the hub monitoring server and portal server immediately disables access to the monitoring data used by the portal client and monitoring applications in Dashboard Application Services Hub.

Locating the portal server in the same LAN segment as the hub monitoring server provides optimal performance when using the Tivoli Enterprise Portal graphical user interface (GUI) and monitoring dashboards in IBM Dashboard Application Services Hub. For users with multiple data centers, if the network connectivity is good, then one portal server is sufficient. If the network connection between data centers has significant latency, then additional portal servers can be placed at each data center.

Caution must be taken when using this approach because there can be only one read-write master portal server. Other portal servers must be used in a read-only manner. Read-only means that users must *not* create and edit objects like situations, workspaces, and policies. Customization must be replicated from the master portal server to the read-only portal server.

If you are using Dashboard Application Services Hub in an environment with read-write and read-only portal servers, you can only use one of the portal servers to retrieve data for monitoring dashboards. Enable the dashboard data provider in one of the portal servers and configure a data provider connection from Dashboard Application Services Hub to this portal server. Because the dashboards only retrieve data and do not make updates to the portal server database, it does not matter which portal server you use to enable the data provider.

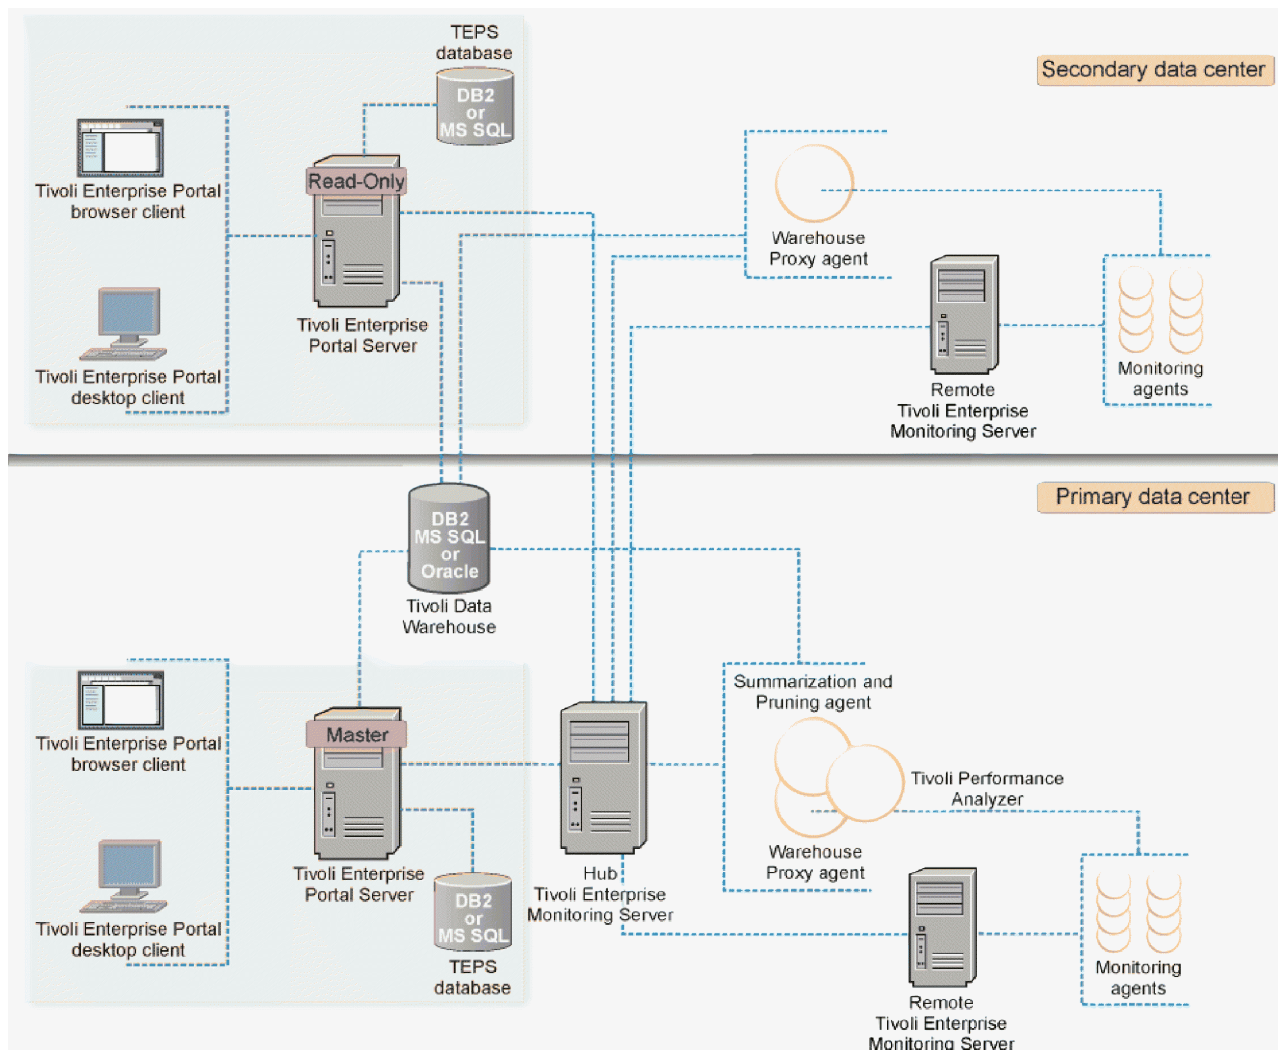


Figure 6. Multiple data center environment

Tivoli Enterprise Portal client

The Tivoli Enterprise Portal client, also known as the portal client, is a Java-based user interface that connects to the portal server to view all monitoring data collections. The portal client is the user interaction component of the presentation layer.

The portal client brings all of these views together in a single window so you can see when any component is not working as expected. The client offers three modes of operation, that all work with Java:

- Browser client
- Desktop client
- Java Web Start client

See Table 3 on page 40 for more information.

The browser client is automatically installed with the Tivoli Enterprise Portal Server (on the Web server that is integrated with the portal server). The desktop client is supported on Windows and Linux operating systems. You can install the desktop client from the IBM Tivoli Monitoring installation media or

you can use IBM Web Start for Java to download the desktop client from the Tivoli Enterprise Portal Server. To find out what browsers are supported see “Supported browser versions” on page 90.

Using browser mode or Web Start clients allow you to perform maintenance updates in a single location. If the desktop client is installed from installation media, maintenance must be installed on each computer. If Web Start for Java is used to download and run the desktop client, you gain the performance advantage of the desktop client along with the convenience of centralized administration from the server. Additional performance gains can be made by modifying the Java heap settings. (For more details on the heap settings see “Locating and sizing the portal client” on page 40.) Unless you want a very secure environment where there are no downloads, use IBM Web Start for Java for obtaining the desktop client.

Note: To use Java Web Start to download the desktop client from the Tivoli Enterprise Portal Server, either Java, version 6.0 or 7.0 (also referred to as JRE version 1.6 or 1.7) must be installed on the system to which you are downloading the client. You can use either the IBM Runtime Environment for Java and the Oracle Java Runtime Environment. The IBM Runtime Environment for Java can be downloaded from the Tivoli Enterprise Portal Server to the system where you are downloading the client. For more details, see “Using Web Start to download and run the desktop client” on page 367.

Many customers install the desktop client on Citrix for the following reasons. Citrix allows for better GUI performance for users at remote locations. In addition, some users do not allow the Tivoli Enterprise Portal client's prerequisite Java version to be installed on desktop computers. By using Citrix, users do not have to install Java on the user's desktop systems.

tacmd Command Line Interface

The Tivoli administrative commands (tacmd) command line interface (CLI) is used to perform administrative tasks such as creating and managing a situation or starting and stopping an agent. It can be installed on Windows and UNIX based operating systems with other IBM Tivoli Monitoring components or on an administrator's computer system.

Some of the tacmd commands send requests to the Hub monitoring server and other commands send requests to the portal server. Therefore, the tacmd CLI needs network access to the Hub Tivoli Enterprise Monitoring Server and to the Tivoli Enterprise Portal Server using the protocols and ports shown in Figure 4 on page 22.

The tacmd CLI users also must specify a user ID and password that the Hub monitoring server or the portal server can authenticate depending on which command is issued.

Tivoli Enterprise Monitoring Agents

Monitoring agents are installed on the system or subsystem that requires data collection and monitoring. Monitoring agents are responsible for gathering data on various properties, or *attributes*, of a monitored system, subsystem, application, or database, the *managed system*, and for sending that data to the monitoring servers.

Monitoring agents test for specified conditions, or *situations*, by periodically comparing attribute values against specified thresholds. When the tested values match or exceed the thresholds, monitoring agents notify the monitoring server, and alert are displayed in the portal client.

The following scenarios prompt the monitoring server to gather data samples from the agents:

- Opening or refreshing a workspace that has data views (table or chart views):
When a Tivoli Enterprise Portal workspace is opened or refreshed, the portal server sends a sampling request to the hub monitoring server. The request is passed directly to the monitoring agent, if there is a direct connection, or indirectly through the remote monitoring server to which the monitoring agent connects. The monitoring agent takes a data sampling and returns the results through the monitoring server and portal server to the portal workspace.
- The sampling interval for a situation (a test taken at your monitored systems):

A situation can have a sampling interval as frequent as once every 30 seconds or as seldom as once every three months. When the interval expires, data samples are requested from the agent in which the returned values are compared with the condition described in the situation. If the values meet the condition, an event is generated and corresponding alerts or automation are initiated.

- Monitoring agents are configured to collect historical data:

The collection of historical data can be configured to send the data to the remote monitoring server at regular intervals or, transfer data collections from the agent to the Warehouse Proxy Agent at hourly or daily intervals. If firewall restrictions are disabled or minimal, configure all of the agents to transfer directly to Warehouse Proxy Agent. If firewall security is an issue, you can either use the firewall gateway or place the Warehouse Proxy Agent on the remote monitoring server where a network connection is already established.

Typically, there are no unique considerations regarding the placement of monitoring agents. The main consideration is whether to store historical data on the agent or on the remote monitoring server. For environments with slow network connections between the agent and the remote monitoring server, storing the data on the monitoring server spreads out the network load of transmitting the historical data. Doing this places a higher demand on the remote monitoring server, and effectively lowers the number of agents that can be managed from the remote monitoring server.

For short-term historical requests of agent data less than 24 hours, the request does not have to cross the slow link, but this savings is offset by the remote monitoring server having to read a larger amount of data from disk (from all agents of that agent type connected to the remote monitoring server) to find the results for the desired agent. For environments with network events that cause agents to failover to a secondary monitoring server, this option may be a poor choice. When an agent fails over, it loses access to the short-term historical data because the data is located on the primary remote monitoring server and the agent is connected to the backup remote monitoring server.

By installing the operating system agent on the system first, the remaining agents can be deployed remotely using the *Add* agent capabilities.

Warehouse Proxy Agent

The Warehouse Proxy Agent is a unique agent that performs only one task: collecting and consolidating all historical data from the individual agents to store in the Tivoli Data Warehouse. If you are using the Tivoli Data Warehouse, at least one Warehouse Proxy Agent is required for each Tivoli Monitoring installation.

The Warehouse Proxy Agent uses ODBC (open database connectivity) on Windows and JDBC (Java Database Connectivity) on AIX and Linux to write the historical data to a supported relational database. The Warehouse Proxy Agent is typically placed on the same LAN segment as the warehouse database, which allows for the best throughput to the database. You can also place the Warehouse Proxy Agent on the same server as the warehouse database.

For larger environments or deployments with multiple data centers, use multiple Warehouse Proxy Agents. You can have as many as one Warehouse Proxy Agent per monitoring server in your monitoring environment. Placing the Warehouse Proxy Agent on the same server as the remote monitoring server reduces the number of servers running Tivoli Monitoring components. This placement also simplifies the configuration because the Warehouse Proxy Agent can be configured to service agents connected to the local monitoring server. If the server running the remote monitoring server goes down and the agent fails over to a secondary monitoring server, the agent should be able to upload historical data through the Warehouse Proxy Agent on the secondary monitoring server. This placement also has the advantage of limiting the number of agents that upload historical data through a single Warehouse Proxy Agent.

NAT environments require specific considerations. The agents receive the IP address of the Warehouse Proxy Agent from the monitoring server. In a NAT network, you must ensure that the agents receive the correct IP address of the Warehouse Proxy Agent. To ensure that the agents can communicate with the

Warehouse Proxy Agent, place the Warehouse Proxy Agent on the same servers as the remote monitoring server. This placement is required only for the agents that are connected to a NAT network.

For information on configuring the historical collection, see the *IBM Tivoli Monitoring Tivoli Enterprise Portal User's Guide*.

Warehouse Summarization and Pruning Agent

The Summarization and Pruning Agent is a unique agent that performs the aggregation and pruning functions for the historical detailed data on the Tivoli Data Warehouse. The Summarization and Pruning Agent has advanced configuration options that enable customization of the historical data storage. One Summarization and Pruning Agent manages the historical data in the Tivoli Data Warehouse.

The Summarization and Pruning Agent can be placed on the Tivoli Data Warehouse database server to minimize network transmission delay during its processing. If the Summarization and Pruning Agent is placed on a separate server, ensure that it connects to the database server with a high-speed network connection of 100 Megabits per second or higher. For large environments, the database server must have at least four processors and a large number of disks, with maintenance by a skilled database administrator.

Performance enhancements were added into the Summarization and Pruning Agent for the IBM Tivoli Monitoring V6.2.1 release. Some of the enhancements require a new database schema. The warehouse will function without the schema changes, but will not take advantage of some of the performance improvements. For the enhancements that do not require a new schema, there are minor database changes that are documented in Chapter 20, "Schema Publication Tool," on page 525. In a new Tivoli Monitoring environment, the warehouse database schema will be created using the new schema. If the Tivoli Monitoring V6.3 environment was updated from a release that used the old schema, the database will continue to use the old schema and will not benefit from the improvements. If you want to take advantage of the performance improvements in an upgraded environment, you will need to create a new warehouse database with a new schema. If desired, you can migrate the data from your old warehouse database into your new warehouse database.

Tivoli Data Warehouse

The Tivoli Data Warehouse is the storage database that contains all of the warehoused (long-term) historical data.

A Warehouse Proxy Agent must be installed to leverage the Tivoli Data Warehouse function within the environment. In large-scale deployments, a Tivoli Data Warehouse can be shared among monitoring installations.

Note: If installing the warehouse database on a Microsoft SQL Server, you must also install the Tivoli Enterprise Portal Server on a Windows-based computer. This restriction applies even if the warehouse database and portal server are installed on separate computers. For example, a portal server on Linux does not support a warehouse database using Microsoft SQL Server.

Dashboard Application Services Hub with Tivoli Common Reporting

Dashboard Application Services Hub is a Jazz for Service Management web-based systems management integration platform that can be used for reporting services.

The Tivoli Common Reporting component of Jazz for Service Management is installed with Dashboard Application Services Hub and provides the following capabilities:

- Common platform for running report packages provided by Tivoli products including IBM Tivoli Monitoring. Also extends the capability of cross-product report linkage for seamless flow between related reports of various Tivoli products.
- Simplified interfaces and interactions for report writers that enable them to create, customize, and publish reports quickly and easily.

Report packages for monitoring agents and custom monitoring reports use Tivoli Data Warehouse as the data source for monitoring data.

Tivoli Common Reporting, dashboard applications such as IBM Infrastructure Management Dashboards for Servers, and the Authorization Policy Server can also be installed in a single Dashboard Application Services Hub for an environment with a small reporting and dashboard workload.

See *Installing Jazz for Service Management* for additional considerations on setting up your reporting environment: http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html.

Netcool/OMNIBus integration

If you are already using Netcool/OMNIBus to monitor events from other sources in your enterprise, you can also view and manage situation events from a Tivoli Enterprise Monitoring Server in the Netcool/OMNIBus console. Event integration requires Netcool/OMNIBus V7.2.1 or later and IBM Tivoli Netcool/OMNIBus Probe for Tivoli EIF version 10 or later.

Situation events are sent to Netcool/OMNIBus Probe for Tivoli EIF using the Tivoli Event Integration Facility (EIF) interface. The Netcool/OMNIBus EIF Probe receives events, maps them to the Netcool/OMNIBus events format, and then inserts into Netcool/OMNIBus ObjectServer. When an Netcool/OMNIBus user acknowledges, closes, or reopens a situation event, Netcool/OMNIBus sends those changes to the originating monitoring server through the event synchronization component.

For information on forwarding events to Netcool/OMNIBus and installing the event synchronization component, see Chapter 30, “Setting up event forwarding to Netcool/OMNIBus,” on page 697.

By default, the EIF Probe listens on the default port (9998).

Tivoli Enterprise Console integration

Tivoli Enterprise Console events can be forwarded from Tivoli Monitoring to Tivoli Enterprise Console Version 3.9. The events are forwarded from the hub monitoring server to the Tivoli Enterprise Console server or Tivoli Enterprise Console gateway. Make sure that firewall ports are opened between the hub monitoring server and Tivoli Enterprise Console servers. By default, the Tivoli Enterprise Console uses port 5529.

The Tivoli Enterprise Console event synchronization component sends updates to situation events back to the monitoring server and are then forwarded to the portal server. Actions performed at the Tivoli Enterprise Console for Tivoli Monitoring situations are reflected in the Tivoli Enterprise Portal Server, this is an optional component that must be installed on your Tivoli Enterprise Console server.

For information on forwarding events to Tivoli Enterprise Console and installing the event synchronization component, see Chapter 31, “Setting up event forwarding to Tivoli Enterprise Console,” on page 799.

Dashboard Application Services Hub with dashboard applications and Authorization Policy Server

Dashboard Application Services Hub is a Jazz for Service Management web-based systems management integration platform with the following features:

- Support for data processing and rendering of content from multiple sources, which is integrated and presented in interactive dashboards.
- Support for multiple user personas, for example, system integrators, administrators, end users, and application developers.
- Support for single sign-on and a unified security mode.

The Dashboard Application Services Hub sends requests to the IBM Tivoli Monitoring dashboard data provider component in the portal server to retrieve agent data for display in the monitoring dashboards such as the IBM Infrastructure Management Dashboards for Servers or in custom dashboards.

The Authorization Policy Server is installed with Dashboard Application Services Hub and allows you to define roles and permissions, which control the access that dashboard users have to managed systems and managed system groups displayed in monitoring dashboards. There is not a graphical user interface for the Authorization Policy Server. Instead administrators use the `tivcmd` CLI to create and work with roles and permissions that are centrally managed by the Authorization Policy Server. The CLI uses HTTP or HTTPS to communicate with the Authorization Policy Server.

The portal server also sends HTTP or HTTPS requests to the Authorization Policy Server application to retrieve updated authorization policies at a configured interval.

You can install an HTTP server for load balancing across a cluster of Dashboard Application Services Hubs. However, the Authorization Policy Server is not currently supported for load balancing in this release.

If you are using Dashboard Application Services Hub in an IBM Tivoli Monitoring environment with multiple portal servers, you can only use one of the portal servers to retrieve data for monitoring dashboards for this release. Enable the dashboard data provider in one of the portal servers and configure a data provider connection from Dashboard Application Services Hub to this portal server.

See *Installing Jazz for Service Management* for additional considerations on setting up your reporting environment: http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html.

tivcmd Command Line Interface for Authorization Policy

An IBM Tivoli Monitoring administrator uses the `tivcmd` Command Line Interface (CLI) to create and work with roles and permissions that control which users and user groups can view managed systems and managed system groups in monitoring dashboards applications or custom dashboards.

The authorization policies are saved in a master policy file store that is maintained on the system where the Authorization Policy Server is installed with Dashboard Application Services Hub. The `tivcmd` Command Line Interface can be located on the same system as the Authorization Policy Server. It can also be installed on an administrator's computer system and uses HTTP or HTTPS to communicate with the Authorization Policy Server.

Administrators who use the `tivcmd` Command Line Interface might also need to use the `tacmd` Command Line Interface to create and work with managed system groups or to perform other administrative functions for IBM Tivoli Monitoring. Both CLIs can be installed on the same system.

Tivoli Enterprise Monitoring Automation Server, Performance Monitoring service provider, Registry Services, and Security Services

The Tivoli Enterprise Monitoring Automation Server is installed with the Hub monitoring server on Windows and UNIX based operating systems. It extends the Hub monitoring server by providing the Open Services Lifecycle Collaboration Performance Monitoring (OSLC-PM) service provider. The service provider registers monitoring resources such as computer systems, software servers, and databases with the Registry Services and also responds to HTTP GET requests for resource health metrics from OSLC clients.

Registry Services is a Jazz for Service Management component that is installed into WebSphere Application Server. Products that discover and manage shared IT resources can register these IT resources and the services they offer with Registry Services. Other products (OSLC clients) can consume data by

querying Registry Services for the managed resources or the associated service providers of interest. In the query responses, Registry Services returns links to the service providers that have more information about the managed resources.

Security Services is an optional Jazz for Service Management component that enables non-WebSphere based applications such as the Performance Monitoring service provider to participate in LTPA based single sign-on. It should be installed and configured if you want the Performance Monitoring service provider to authenticate requests it receives from OSLC clients. This is the only authentication method supported by the service provider. The application servers for OSLC clients also must be configured for single sign-on with Security Services and include an LTPA token in HTTP GET requests sent to the service provider. When the Performance Monitoring service provider receives an HTTP GET request from an OSLC client, it forwards the LTPA token to Security Services to authenticate the request. If the request does not contain an LTPA token or Security Services indicates that the token is not valid or has expired, the Performance Monitoring service provider returns an HTTP 401 status code to indicate that the request could not be authenticated.

For the current release of IBM Tivoli Monitoring, the Performance Monitoring service provider assumes that Registry Services and Security Services are installed into the same WebSphere Application Server.

See *Installing Jazz for Service Management* for recommendations on whether to install Registry Services and Security Services features on the same systems as other Jazz for Service Management components: http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html.

IBM Tivoli Agent Builder

The Agent Builder is a wizard that makes it easier to build custom monitoring solutions. Using mostly point and click capabilities, the Agent Builder allows you to create a monitoring agent using multiple data providers such as WMI, Perfmon, Log scraping, scripts, and process monitoring. Over time additional data providers will be added to the Agent Builder.

For more information about Agent Builder, see the *Agent Builder User's Guide*.

Note to Windows users: Since the introduction of the Embedded Java Runtime and the Tivoli Enterprise Services User Interface Extensions (the KUE component) at IBM Tivoli Monitoring V6.2.2, the possibility has existed that, on nodes where only the Windows OS agent was installed, the Embedded Java Runtime and the User Interface Extensions were not also installed. If you later attempt to install an Agent Builder agent on such nodes, you may receive the error shown in Figure 52 on page 267. If this happens, complete the procedure outlined in “Installing the Embedded Java Runtime and the User Interface Extensions” on page 267, and retry the agent installation.

Firewall gateway

Using the firewall gateway, you can traverse even the most complex firewalls.

Using the IP PIPE or IP SPIPE protocols, the Tivoli Monitoring software can traverse most firewall configurations. In most cases, the firewall gateway is not necessary.

For detailed information on installing and configuring the firewall gateway, go to Appendix C, “Firewalls,” on page 861.

Tivoli Performance Analyzer

Tivoli Performance Analyzer adds predictive capability to Tivoli Monitoring so you can monitor resource consumption trends, anticipate future performance issues, and avoid or resolve problems more quickly. For example, you can use Tivoli Performance Analyzer to predict application bottlenecks and create alerts for potential service threats.

Tivoli Performance Analyzer helps IT managers to answer the following questions so they can optimize IT capacity:

- When will my application fail to meet service levels?
- How will application performance change if I modify the infrastructure?
- What is the best hardware solution to meet my performance and cost goals?
- Where are my under-utilized servers and networks?
- Which servers and network components really need an upgrade?
- Which application will experience the next performance issue? When?

Providing accurate IT forecasts and appropriate IT capacity to meet business needs are two major goals in the capacity management area. You can use the following key performance indicators (KPIs) to measure the critical success factors:

- Total value of unplanned or unused capacity expenditures
- Percent of capacity forecasts that were accurate
- Number of inaccurate business forecast inputs provided
- Number of incidents related to capacity or performance issues

Monitoring agent for IBM Tivoli Monitoring 5.x Endpoint

Also called IBM Tivoli Monitoring 5.x Endpoint Agent, this integration agent enables the collection and visualization of IBM Tivoli Monitoring 5.x resource models in the Tivoli Enterprise Portal.

The visualization is the direct replacement for the Web Health Console. Additionally, the agent provides the roll-up function of IBM Tivoli Monitoring 5.x metrics into the Tivoli Data Warehouse.

Additional ports used in the Tivoli Monitoring environment

If multiple components are installed on the same server, they may not share the same port number for IP.PIPE or IP.SPIPE communications. A scheme has been created to use additional ports for communications.

See the section, “Controlling port number assignments” on page 421 for more information on managing ports when you have multiple Tivoli Monitoring components running on a single system.

Configuring your firewalls

IBM Tivoli Monitoring has a firewall gateway component that allows you to traverse even the most complex firewalls.

Figure 4 on page 22 outlines the ports and communications that must occur between the various Tivoli Monitoring components, and shows an example for a typical environment using IP.PIPE. If you are using IP.SPIPE, the diagram would look similar, but port 3660 would be used instead of 1918. The same skip counts apply as described in “Controlling port number assignments” on page 421, but the initial port number is 3660.

Sizing your Tivoli Monitoring hardware

The following section outlines hardware scenarios for Tivoli Monitoring environments of various sizes.

These are rough guidelines based on typical deployments. Because Tivoli Monitoring is highly customizable, this section also outlines some usage scenarios that drive additional load and might indicate different hardware requirements.

“Required hardware for distributed systems” on page 84 has a section describing the hardware requirements for the Tivoli Monitoring infrastructure components, including the processor, memory and disk requirements. The guidelines below are based on typical deployments, and supplement the hardware requirements described in Chapter 3, “Hardware and software requirements,” on page 75.

Locating and sizing the hub Tivoli Enterprise Monitoring Server

Always locate the hub monitoring server in a data center with good network reliability and throughput. Network connection speeds of 100 Megabits per second or higher are typically sufficient.

Except where noted, the server sizings assume that the hub monitoring server is installed on a separate server with no additional Tivoli components. Below are some basic guidelines. For detailed guidelines, see Chapter 3, “Hardware and software requirements,” on page 75.

Sizing your hub monitoring server hardware

Large and medium-sized environments (>1500 agents)

Remote Tivoli Enterprise Monitoring Servers should be used in this environment. The Hub Tivoli Enterprise Monitoring Servers should have at least 2 processors and 3 GB of RAM.

Although the steady-state CPU utilization is usually low, the hub monitoring server uses a significant amount of system resources in processing transient workloads, such as startup, agent login, and situation distribution. It is reasonable to run other Tivoli Monitoring components on the same server such as the portal server, but allocate additional CPUs for the additional components. For more information, see “Memory and disk requirements” on page 85.

Small environments (< 1000 agents)

Remote Tivoli Enterprise Monitoring Servers are not typically required in this environment. The system should have at least 2 processors and 4 GB of RAM. It is reasonable to combine components such as the monitoring server and portal server on a dual-processor server. When combining monitoring components on a single server, the different components' process memory requirements should be added together. See “Memory and disk requirements” on page 85 for more information.

For environments with 200 or fewer agents, most of the components can be combined onto a single server. If you are going to use a small monitoring environment, choose a multiprocessor (2 or 4-way) for the monitoring server with at least 4 GB of total system memory. Configure the software and begin monitoring CPU and memory usage during periods of both normal and high volumes of situation events. If CPU or memory usage is constrained, consider deploying a separate server for the Tivoli Data Warehouse, the Warehouse Proxy Agent, and the Summarization and Pruning Agent. For more information see “Memory and disk requirements” on page 85.

When to use multiple hub monitoring servers

As environments grow in size, you must run multiple hub monitoring servers. The recommended limit of managed systems per hub monitoring server is 20,000.

There are many factors that affect the scalability of a hub environment. The primary factors are network latency, number of situations running at the hub, the number of concurrent active users, and historical data collection.

Note: When multiple hub monitoring servers are used, correlation of events can be performed by using either Tivoli Enterprise Console or Netcool/OMNIBus.

Locating and sizing the remote Tivoli Enterprise Monitoring Server

When determining the location of your remote monitoring server, you must consider the number of connected agents, existence, number, and placement of firewalls, and network bandwidth. With the Tivoli Monitoring V6.3 release, as a general planning guideline, a single remote monitoring server can support up to 1500 managed systems.

The main factor affecting remote monitoring server scalability is situation processing, and the volume and frequency of agent attribute group data that is processed by the monitoring server during situation

evaluation. Situation processing affects monitoring server performance, so the CPU and memory usage of the kdsmain process should be monitored. For tuning recommendations for improving situation processing performance, see “Optimizing situations” on page 504.

Although the steady state CPU utilization is usually low, a remote monitoring server uses a significant amount of system resources in processing transient workloads, such as startup, agent login, and situation distribution. Dual processor systems are recommended for remote monitoring servers. To support 1500 managed systems, a server with 4 GB of memory is typically sufficient. For environments with more than 3000 agents collecting historical data, you can use multiple Warehouse Proxy agents, with the Warehouse Proxy agents located on each remote monitoring server to handle historical data exports from agents connected to the remote monitoring server.

If the remote deployment capability is used, ensure that the remote monitoring server has sufficient network capabilities (100 Megabits per second or higher is typically sufficient). Because the remote deployment packages are large, in some cases hundreds of megabytes, you must ensure that you do not have a network bottleneck.

The network connection speed between the monitoring server and the monitoring agent will affect the elapsed time for remote deployment operations. One way to deal with slow network links is to deploy a remote monitoring server near the monitored servers. By co-locating the remote monitoring server with the agents, the Remote Deploy bundles need to be transferred across the slow link only once. Distribution to the agents can then be done on a high speed LAN segment.

Locating and sizing the remote deployment depot

When you use the remote deployment for agents and patches, place the remote deployment depot on a separate directory from your Tivoli Monitoring installation directory. This allows for cleaner backup and restore operations.

Note: Remote deployment is possible only for monitoring agents that run on distributed platforms (Windows, UNIX, and Linux computers).

The remote deployment depot requires significant disk space, which varies depending on the number of agents and the number of platforms you use. Typically, allocate at least 2 GB of disk space for the depot.

Note: If you create a shared deployment repository named depot on the server hosting the deployment depot and you create this repository in a subdirectory of the depot directory, the monitoring server will not be able to find your deployment depot. Instead:

1. Create the repository at the C:\IBM\ITM\CMS level of the directory structure, *not* at the C:\IBM\ITM\CMS\depot level.
2. Then set the DEPOTHOME keyword to DEPOTHOME=\\hubtems\centralrepository\depot in the KBBENV file.

Otherwise you will get this error message:

```
KDY2077E: The specified agent bundle depot \\hubtems\depot is not a directory.  
Either the agent bundle depot directory does not exist or it is not a directory.  
The agent bundle depot directory does not exist because no bundles have been added.
```

One of the most important aspects of using remote deployment is ensuring that the files that are distributed are at the correct version. In large environments, there might be several remote monitoring servers that require maintenance to ensure that all of the depots are updated. To avoid the necessity of installing updates and maintenance on multiple computers, consider creating one depot and making it available using NFS or Windows shares.

If you decide to use a shared depot, make sure that the remote monitoring server is upgraded before you upgrade the agents that are connected to that remote monitoring server. When you deploy monitoring agents, make sure that you configure them to connect to the desired remote monitoring server.

If you are using a Windows monitoring server connecting to a depot on another Windows computer, you must set the service ID for the Windows monitoring server to "Administrator." Also, instead of specifying a mapped drive letter for the path to the depot directory, use the UNC path (such as \\server\share). Use the following steps to change the service ID:

1. From the Control Panel, double-click **Administrative Tools**.
2. Double-click **Services**.
3. Right-click **Tivoli Enterprise Monitoring Svcs** and click **Properties**.
4. On the **Log On** tab, select **This Account**.
5. Type Administrator in the **This Account** field.
6. Type the password for the administrator in the **Password** field. Confirm the password by typing it again in the **Confirm password** field.
7. Click **Enable**.

If the Administrator user does not have Logon as a service right, you are prompted to add it.

Note:

1. In addition to remotely deploying monitoring agents, IBM Tivoli Monitoring lets you remotely deploy *non-agent bundles* (with which your site can employ Tivoli Monitoring components that need not connect to a Tivoli Enterprise Monitoring Server).
2. Remote deployment is not available for z/OS-based OMEGAMON agents, nor is it supported on nodes running the Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, or the Tivoli Enterprise Portal desktop or browser client.

Locating and sizing the Tivoli Enterprise Portal Server

Install the portal server on a very fast server with at least 2 processors. Memory requirements are described in Table 18 on page 85.

The Tivoli Monitoring V6.3 release has been tested with 100 concurrent portal clients on 64-bit platforms. See the "Tivoli Enterprise Portal Server" on page 474 for more details.

For optimal performance, place the portal server in the same data centers as most of the end users. Depending on firewall restrictions the location of the portal server has to be verified. If clients are not allowed to connect to a central segment they are not able to use the portal server.

If you have multiple large data centers you might need to install a second read-only portal server. If you choose to use this option, be aware that only one portal server can be used to make customization changes such as editing workspaces and situations. See the section on high availability for details on how to configure multiple portal servers.

For more information see "Memory and disk requirements" on page 85.

Locating and sizing the Warehouse Proxy Agent

The Warehouse Proxy Agent is driven by the amount of historical data being collected. The Warehouse load projection spreadsheet can help you estimate both the number of database inserts and the volume of data being collected in the warehouse database.

You can find the Warehouse load projection spreadsheet in the Tivoli Integrated Service Management Library by searching for "warehouse load projections" or the navigation code "1TW10TM1Y" at IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>).




Figure 7. Warehouse load projection spreadsheet

Figure 7 shows a sample spreadsheet summary created using the tool downloaded from the Tivoli Integrated Service Management Library site. The amount of network traffic closely matches the amount of data collected in the Warehouse. The key planning numbers, highlighted in red, are based on sample data.

Each Tivoli Monitoring environment is different, so fill out the spreadsheet based on your warehousing needs. The key information in the spreadsheet includes the following data:

- Total Tivoli Data Warehouse record inserts per hour
In this example, there are 139,288 inserts per hour. Most reasonable database servers can handle this rate without any problem.
- Total megabytes of detailed data inserted per hour
In the example, 81.6 MB of data per hour are inserted. This amount roughly translates to 81.6 MB of network traffic per hour coming into the Warehouse Proxy Agent from the connected agents, and roughly 81.6 MB per hour of network traffic from the Warehouse Proxy Agent to the warehouse database server.
- Total gigabytes of data in Tivoli Data Warehouse (based on retention settings)
In this example, you expect the Tivoli Data Warehouse database to grow to 49.4 GB of data. Additional space must be allocated for the database logs. The database does not grow beyond 49.4 GB because you are summarizing the raw (detailed) metrics and pruning both the detailed and summarized data. The planning spreadsheet helps you determine your warehouse size based on your retention needs.

Multiple Warehouse Proxy Agents are required when the number of agents collecting historical data increases above approximately 3000 (to ensure that limits for the number of concurrent connections are

not reached). If you use multiple Warehouse Proxy Agents, consider running them on the same servers running the remote monitoring servers, and configuring them to support agents connected to this monitoring server. This approach consolidates the Tivoli Monitoring infrastructure components, and limits the number of agents that connect to each Warehouse Proxy Agent.

For a single Warehouse Proxy agent supporting up to 3000 agents, a dual-processor system with 2 GB RAM should be sufficient. For running multiple Warehouse Proxy agents, they are typically located on the remote monitoring servers.

Locating and sizing the Summarization and Pruning Agent

The Summarization and Pruning Agent provides the ability to customize the length of time for which to save data (pruning) and how often to aggregate detailed data (summarization) in the Tivoli Data Warehouse database. Configuration of the Summarization and Pruning Agent can have a significant impact on the performance of your Warehouse.

The Summarization and Pruning agent can be installed on the warehouse database server, or on a separate fast server with a low latency and high network bandwidth connection to the warehouse database server (100 Mbps or higher).

Table 2 shows guidance for setting up the Summarization and Pruning Agent based on the amount of data inserted into the warehouse database per day. The server sizing is based on the warehouse database and Summarization and Pruning Agent running on the same server, and the Warehouse Proxy Agent running on a separate server. The assumption is that you have new and fast server-class hardware because the warehouse database server is a high-performing server.

For smaller environments, use a minimum of three dedicated disks because of the high disk I/O requirements of the warehouse database server.

Table 2. Warehouse database server considerations

Data inserted per day	Number of CPUs	Memory	Hard disk storage
0 to 500 K inserts per day	2 CPUs	2 GB	4 or more dedicated SCSI drives
0.5 to 2 Million inserts per day	2 CPUs	4 GB	RAID Array with 4 or more disks
2 to 10 Million inserts per day	4 CPUs	4 GB	Multiple RAID Arrays with 5 or more disks. Or, SAN storage
10 to 20 Million inserts per day	4 CPUs	8 GB	RAID Arrays with 15 to 20 dedicated disk drives or high-performance SAN storage
> 20 Million inserts per day	4 CPUs	8 - 16 GB	Multiple RAID Arrays with 20 to 25 dedicated disks or high-performance SAN storage

If your Summarization and Pruning Agent is installed on the warehouse database server, configure the agent so that it leaves some processing power for others to perform Warehouse queries and for the Warehouse Proxy Agent to insert data. Set the number of worker threads to 2 to 4 times the number of CPUs on the system where the Summarization and Pruning Agent is running. Start with 2 times the number of CPUs and then increase the number to see if it continues to improve your performance.

To set the number of worker threads, edit the Summarization and Pruning Agent configuration file (KSYENV on Windows, sy.ini on UNIX or Linux) and set the **KSY_MAX_WORKER_THREADS** to the number of desired threads. This parameter can also be set using the configuration dialog panels.

To assess the performance and throughput of the Summarization and Pruning Agent for your environment you can use the Statistics workspace and the Work Completed workspace. These workspaces are described in the *Warehouse Summarization and Pruning Agent User's Guide*. See “Tivoli Data Warehouse” on page 480 for tuning information.

Locating and sizing the portal client

There are three deployment options for the portal client; Browser client, Desktop client, and Java Web Start client.

The portal client has three different deployment alternatives which are described in Table 3 below.

Table 3. Portal client deployment considerations

Portal client	Advantages	Disadvantages
Browser client	<ul style="list-style-type: none">• Does not need to be installed on client machine.• Centralized administration.• Workspaces can be referenced using URLs.	<ul style="list-style-type: none">• Longer initial download.
Desktop client	<ul style="list-style-type: none">• Fastest startup and performance.	<ul style="list-style-type: none">• Needs to be installed on each client machine.• Requires maintenance to be installed individually on each client machine.• Mismatch of versions between portal server and the client is not permitted.
Java Web Start client	<ul style="list-style-type: none">• Similar performance to desktop client.• Supports multiple JREs.• Centralized administration.	<ul style="list-style-type: none">• Longer initial download.

The memory required by the portal client depends on the size of the monitoring environment and on the Java heap size parameters. The default maximum Java heap size for the desktop client is 256 MB on 32-bit platforms and 512 MB on 64-bit platforms. For a large monitoring environment, the portal client can be expected to use between 300 and 600 MB of memory.

Locating and sizing the Tivoli Performance Analyzer

The Tivoli Performance Analyzer can be installed onto a stand-alone workstation, onto the same workstation as the Tivoli Monitoring environment if it is installed on a single machine, or onto the same workstation as any other Tivoli Monitoring element in a distributed environment.

For large environments, it might be necessary to increase the maximum Java heap size to 512 MB. For more information, see “Setting the maximum Java heap size” on page 487. Before installing Analytic Agent make sure that your Tivoli Monitoring environment includes the listed products, that they are configured and that you have administrator access to the following components:

- Tivoli Data Warehouse
- Tivoli Enterprise Monitoring Server
- Tivoli Enterprise Portal Server
- Summarization and Pruning Agent

Note: Tivoli Performance Analyzer analyzes summarized data. Therefore, the Summarization and Pruning Agent must be installed and active so that data is available for the agent to analyze.

- Warehouse Proxy Agent

Note: The Performance Analyzer agent application support that is installed in the monitoring server, portal server, and portal client must be at the same version as the Performance Monitoring agent.

The analytical calculations performed by Tivoli Performance Analyzer use the data in Tivoli Data Warehouse. Your Tivoli Monitoring installation must therefore be configured for History Collection and Warehouse Summarization and Pruning.

Elements to install

Tivoli Performance Analyzer requires other elements to be installed.

There are several distinct elements that must be installed to add Tivoli Performance Analyzer into your Tivoli Monitoring environment:

- Performance Analyzer Agent
- Tivoli Enterprise Monitoring Server Support
- Tivoli Enterprise Portal Server Support
- Tivoli Enterprise Portal Support

Note: The Performance Analyzer agent application support that is installed in the monitoring server, portal server, and portal client must be at the same version or higher than the Performance Monitoring agent.

Support files must be installed on the same workstation as the components they are supporting. For example, Tivoli Enterprise Monitoring Server support should be installed on the same machine as Tivoli Enterprise Monitoring Server.

For information on installing the domain support to enable performance analytics for a broader set of systems, see “Installing domain definitions for Tivoli Performance Analyzer” on page 336.

Deployment scenario

Example distributed installation of Tivoli Performance Analyzer and associated support files.

Table 4 shows an example distributed installation of Tivoli Performance Analyzer. The agent is installed on the same workstation as Tivoli Data Warehouse.

Table 4. Example distributed installation

Machine 1	Machine 2	Machine 3
Tivoli Enterprise Monitoring Server	Tivoli Enterprise Portal Desktop Client	Analytic Agent
Tivoli Enterprise Portal Server		IBM Tivoli Data Warehouse

Table 5 shows the support files that must be installed on each workstation in the example distributed installation shown in Table 4.

Table 5. Support files to be installed

Machine 1	Machine 2	Machine 3
Tivoli Enterprise Monitoring Server Support	Tivoli Enterprise Portal Desktop Client Support	
Tivoli Enterprise Portal Server Support		

Table 5. Support files to be installed (continued)

Machine 1	Machine 2	Machine 3
Tivoli Enterprise Portal Browser Client Support		

Using IBM Installation Manager

IBM Installation Manager is an installation program that uses remote or local software repositories to install, modify, or update certain IBM products. IBM Installation Manager determines and displays available packages, checks prerequisites and interdependencies, and installs or modifies the selected packages. You also use IBM Installation Manager to uninstall the packages that it installed.

A repository is a location that stores data for installing, modifying, rolling back, updating, or uninstalling packages. A repository can be a DVD, an expanded compressed file, a network file server, or an enterprise web server. IBM Installation Manager can also connect to repositories hosted on <http://www.ibm.com> for installing fixes and service packs (provided an internet connection is available).

The *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image includes the IBM Tivoli Monitoring installation launchpad application. The installation launchpad provides a single location from which to start the installation process, and simplifies the installation process by installing IBM Installation Manager on your computer, if it is not already installed. The installation launchpad then automatically configures IBM Installation Manager to locate the installation files for the software that you are installing. This approach only allows you to install products from the repository located on the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image. When running from the launchpad, IBM Installation Manager reports prerequisites and interdependencies, but does not automatically install the prerequisite software. For more information, see “Required software and memory requirements for a dashboard environment” on page 91.

Another common configuration is to combine the repositories of certain IBM products into an enterprise repository. The offerings on the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image coexist with other IBM Tivoli Monitoring products. For example, IBM Tivoli Monitoring for Virtual Environments or components from Jazz for Service Management, allow for easy installation on to multiple servers without having to move disks between machines. When running from an enterprise repository, IBM Installation Manager reports prerequisites and interdependencies, and automatically installs the required software (provided the offerings are loaded in to the enterprise repository).

For instructions on installing your product in an enterprise environment and for more information about IBM Installation Manager, go to <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.

The following links provide detailed information about how to use a repository to install, update, modify, and uninstall packages on your computer using IBM Installation Manager:

- Repository Preferences (http://publib.boulder.ibm.com/infocenter/install/v1r6/index.jsp?topic=/com.ibm.cic.agent.ui.doc/topics/t_specifying_repository_locations.html)
- Installing updates (http://publib.boulder.ibm.com/infocenter/install/v1r6/index.jsp?topic=/com.ibm.cic.agent.ui.doc/topics/t_install_updates.html)
- Modifying installed packages (http://publib.boulder.ibm.com/infocenter/install/v1r6/index.jsp?topic=/com.ibm.cic.agent.ui.doc/topics/t_modify_install_pkg.html)
- Rolling back updates (http://publib.boulder.ibm.com/infocenter/install/v1r6/index.jsp?topic=/com.ibm.cic.agent.ui.doc/topics/t_rollback_updates.html)

Software Product Compatibility Reports (SPCR)

When planning an installation, it is important to review the list of supported platforms to ensure that you are running on a supported operating system and database version.

The Software Product Compatibility Reports that contain the current list of supported platforms can be found at Software Product Compatibility Reports (<http://pic.dhe.ibm.com/infocenter/prodguid/v1r0/clarity/index.html>).

Configuring for high availability and disaster recovery

Among the most important considerations in setting up your Tivoli Monitoring environment is ensuring high availability of the product components and being able to recover quickly from failures.

There are multiple components to consider when discussing high availability. In the following sections, each of the Tivoli Monitoring components is discussed, along with a strategy for achieving the desired level of high availability and disaster recovery. Ensuring high availability involves achieving redundancy for every Monitoring component. Disaster recovery means being able to recover from a major outage such as a data center going offline or losing its WAN link.

Configuring the hub monitoring server for high availability and disaster recovery

The hub monitoring server is a highly reliable component, and many users choose to run with a single hub monitoring server and use backup and restore operations to ensure that they have minimum downtime in case of a hardware failure. Other users require higher availability and less downtime and employ multiple hub monitoring servers to achieve either a high availability (HA) environment, disaster recovery (DR) environment, or a combination (high availability and disaster recovery) environment. The following section describes some of the strategies used to achieve the desired level of availability and downtime.

If you have a smaller environment and do not want to invest in additional hardware, you can set up a single hub monitoring server. Because the hub monitoring server is very reliable, there is no need to purchase any additional hardware. You can expect some downtime when patching the hub monitoring server. There are times when the monitoring server must be recycled. If you use one hub, it is important that you use a good backup and restore strategy. You can install the hub in a virtualized environment such as VMWare so that you can quickly bring up an identical virtual operating system to replace the original. In addition, there is a VMWare HA option with release 3.0.x that automates the start of a failing image on a different node.

If you want to achieve high availability, you have two options. The first option is to implement the Hot Standby feature that is built into the monitoring server. Extensive large scale testing has taken place to ensure that Hot Standby is a robust solution. The second option is to implement an operating system cluster. Extensive testing has been performed with some operating system clusters. The first two supported clustering options are Windows Cluster and High Availability Cluster Multi-Processing (HACMP™).

The main difference between the clustering options is the scripts to control the automated control of resources within the cluster. In this sense, Tivoli Monitoring is cluster-ready for other clustering solutions.

For detailed steps for configuring Clustering, including a detailed clustering scenario, see the *IBM Tivoli Monitoring High Availability Guide for Distributed Systems*. Also see this guide for instructions on configuring the Hot Standby option.

Note: The Tivoli Enterprise Monitoring Automation Server and its Performance Monitoring service provider cannot be used if the Hot Standby option is used.

For more information about using a high availability z/OS hub monitoring server, see *Configuring the Tivoli Enterprise Monitoring Server on z/OS* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm).

Configuring for portal server high availability and disaster recovery

It is important to have multiple portal servers available in case of a hardware failure. While the hub monitoring server tracks the state of your Tivoli Monitoring environment it is not as critical to ensure that data is synchronized in real-time between multiple portal servers. The primary data that you want to protect is the customization that is stored in the portal server database, such as user-defined workspaces. Because this data does not change frequently, a good backup and restore strategy ensures a highly available environment.

Here are two different ways of achieving both high availability and disaster recovery with the portal server, depending on the amount of hardware you are willing to dedicate to your solution:

OS Cluster

Many users set up an OS Cluster for their portal server. Depending on the clustering software used, the cluster can be set up across a WAN to achieve disaster recovery. For detailed information on setting up the monitoring server and portal server in an OS Cluster, see the *IBM Tivoli Monitoring High Availability Guide for Distributed Systems*.

Cold backup

Some smaller users do not want to dedicate CPU cycles and memory to a live backup portal server. If that is the case in your environment, install a second portal server on another computer that serves as a production server. The backup portal server is typically shut down so that it does not use any CPU or memory. If the primary portal server goes down, the cold backup can be brought online. The key for a cold backup portal server is to periodically export the portal server database content and import it into the cold backup. In addition, ensure that the cold backup portal server is patched with the same software levels as the primary portal server and the same application support.

Consider using a tool like the Tivoli System Automation for Multiplatforms to automate the process of backing up the resources.

As discussed previously, some users choose to implement a master read-write portal server and one or more read-only portal servers. When you implement multiple read-only portal servers, you can place a load balancer or edge server in front of the portal server and have users connect to the edge server. By doing this, you minimize the complexity and maximize the availability for the end user.

Using a load balancer between IBM Dashboard Application Services Hub and multiple portal servers is not supported in this release. Enable the dashboard data provider in one of the portal servers and configure a data provider connection from Dashboard Application Services Hub to this portal server. Because the dashboards only retrieve data and do not make updates to the portal server database, it does not matter which portal server you use to enable the data provider.

The strategy for backup and restore is to have one master Tivoli Enterprise Portal Server database where all customization is done. Then, periodically export the content from the “master” portal server and import the content into any other portal server. The import replaces the Tivoli Monitoring content in the portal server database, so be aware that any customization made in the secondary portal server environments will be overwritten during the import. The export and import of the portal server database can be done in two ways:

- Using RDBMS backup utilities such as DB2’s db2 backup and db2 restore commands
- Using the migrate-export and migrate-import command provided by the Tivoli Monitoring product

If the various portal server databases are not running on the same OS version, then the RDBMS backup and restore utilities will probably not work. In those cases, use the Tivoli Monitoring migrate-export and migrate-import commands as described in the product documentation.

Configuring for IBM Dashboard Application Services Hub high availability and disaster recovery

IBM Dashboard Application Services Hub Server supports load balancing for high availability.

See the Jazz for Service Management Configuration Guide for more details on setting up this support, Jazz for Service Management Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/psc_ic-homepage.html).

However, the Tivoli Authorization Policy Server is not supported for load balancing with Dashboard Application Services Hub in this release.

Configuring for agent and remote monitoring server high availability and disaster recovery

All agents can be defined with a primary and secondary monitoring server, which allows the agent to connect to the secondary monitoring server if the primary is unavailable. Failover to the secondary monitoring server occurs automatically if the agent fails to communicate with the primary monitoring server.

If no other communication occurs between the agent and the monitoring server, the longest interval it should take for the failover to occur is the heartbeat interval, which defaults to 10 minutes.

The primary concern when building a high availability and disaster recovery configuration for the agents and remote monitoring servers is to determine how many agents to connect to each remote monitoring server. For Tivoli Monitoring V6.3, no more than 1500 monitoring agents should connect to each remote monitoring server.

The following information is important when planning your agents and remote monitoring servers:

- Ensure that failover does not result in many more than 1500 monitoring agents reporting to a single remote monitoring server. There are two strategies users typically take to avoid this situation.
 - The *first* and preferred strategy involves having a spare remote monitoring server. By default, the spare remote monitoring server has no agents connected. When the monitoring agents that report to the primary monitoring server are configured, they are configured to use the spare remote monitoring server for their secondary monitoring server. Over time, network and server anomalies cause the agents to migrate.

To manage this environment, write a situation to monitor how many agents are connect to the spare remote monitoring server. You can then use the situation to trigger a Take Action command that forces the agents back to their primary remote monitoring server by restarting them. Restarting the agents cause them to connect to their primary monitoring server. Ideally, migrate the agents back to their primary remote monitoring server when the number of agents connect to the spare monitoring server is greater than 20.

The disadvantage to using a spare remote monitoring server is that you must dedicate a spare server to be the spare remote monitoring server. Some users choose to co-locate this server with the Warehouse Proxy Agent or run in a virtualized environment to minimize the extra hardware required.

- The *second* strategy is to evenly distribute the agents so that they failover to different remote monitoring servers to ensure that no remote monitoring server becomes overloaded. In the example below, there are four remote monitoring servers. In this example, configure one-third of the agents on each remote monitoring server to failover to a different remote monitoring server. Review the following scenario:

RTEMS_1 has 1125 agents, RTEMS_2 has 1125 agents, RTEMS_3 and RTEMS_4 have 1125 agents.

A third of RTEMS_1's agents failover to RTEMS_2, a third failover to RTEMS_3, and a third failover to RTEMS_4.

This strategy ensures that none of the remote monitoring servers become overloaded. The problem with this strategy is that it requires a lot of planning and tracking to ensure that all of the remote monitoring servers are well-balanced.

- If you want your agent to failover to a remote monitoring server in another data center, ensure that you have good network throughput and low latency between the data centers.

Note: Connect a very small number of agents to the hub monitoring server. Typically, only the Warehouse Proxy Agent, Summarization and Pruning Agent, and any OS agents that are monitoring the monitoring server are connected to the hub monitoring server.

Use the Tivoli Monitoring heartbeat capabilities to ensure that agents are running and accessible. The default heartbeat interval is 10 minutes. If an agent does not contact the monitoring server, a status of `MS_Offline` is seen at the monitoring server. An event can be generated when an agent goes offline. An administrator can evaluate whether the agent is having problems or whether there is another root cause. In addition, there is a solution posted on the Tivoli Integrated Service Management Library Web site that leverages the `MS_Offline` status and attempts to ping the server to determine if the server is down or whether the agent is offline. You can find more information by searching for "Perl Ping Monitoring Solution" or navigation code "1TW10TM0F" in the IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>).

Configuring for warehouse high availability and disaster recovery

When setting up the Warehouse for high availability and disaster recovery, the primary concern is backing up the data.

The warehouse database can grow rapidly and has significant change, with many gigabytes of new data inserted per day plus summarization and pruning. Use the native database replication tools to achieve a high availability solution. All of the major database vendors provide data replication tools.

Configuring for Warehouse Proxy Agent high availability and disaster recovery

You need to achieve redundancy with the Warehouse Proxy Agent. Only one Warehouse Proxy Agent can be receiving historical data from a specific agent. You can encounter problems if two Warehouse Proxy Agents are configured to receive historical data from the same agent. To avoid problems, ensure that only one Warehouse Proxy Agent is responsible for collecting the historical data from a remote monitoring server.

To ensure that your Warehouse server performs optimally, ensure that the `WAREHOUSELOG` and `WAREHOUSEAGGREGLOG` tables are pruned on a regular basis.

Note: Beginning with Tivoli Monitoring V6.2.3 the tables `WAREHOUSELOG` and `WAREHOUSEAGGREGLOG` are disabled by default.

Pruning for these tables can be configured by specifying retention intervals in the configuration dialog for the Summarization and Pruning Agent or in the configuration file (`KSYENV` on Windows, `sy.ini` on UNIX or Linux). See "Historical data collection" on page 481 for more details.

Configuring for Summarization and Pruning Agent high availability and disaster recovery

Connect the Summarization and Pruning Agent to the hub monitoring servers. When the Hot Standby option is used, the Summarization and Pruning Agent must be configured with the standby hub as the

secondary monitoring server. However, there are some additional considerations for achieving high availability with the Summarization and Pruning Agent.

Only one Summarization and Pruning Agent may be running against a warehouse database. Thus, it is important to ensure that there is data integrity within the database and that there is no database deadlock between two competing agents. So, by default, only one Summarization and Pruning Agent must be installed and running.

As in the Warehouse Proxy Agent set up, you want to install a second Summarization and Pruning Agent that serves as a cold backup to the primary Summarization and Pruning Agent. By default, the backup Summarization and Pruning Agent is disabled. Write a situation that detects when the primary Summarization and Pruning Agent is down and automatically starts up the backup Summarization and Pruning Agent through a Take Action command.

Care must be taken in writing the Take Action command to ensure that only one Summarization and Pruning Agent is running at any given time. To ensure the two Summarization and Pruning Agents are not running at the same time, perform the following steps:

1. Have the situation trigger only after the second or third missed heartbeat. Occasionally, there are temporary outages triggered by network problems or routine maintenance. You do not want the automated Take Action to occur during a temporary outage.
2. When starting up the backup Summarization and Pruning Agent using a Take Action command, the primary Summarization and Pruning Agent must be disabled so that it does not accidentally restart until an administrator manually corrects the problem.
3. Write up a documented procedure to ensure that only one of the Summarization and Pruning Agents is brought back online following a failover.

Configuring for Tivoli Performance Analyzer high availability and disaster recovery

The Tivoli Performance Analyzer must always connect to the hub monitoring server. When you use the Hot Standby option, you can configure Tivoli Performance Analyzer Agent with the standby hub as the secondary monitoring server.

Since there can only be one Tivoli Performance Analyzer Agent in your Tivoli Monitoring environment (that is one Tivoli Performance Analyzer Agent per hub monitoring server) it is not possible to setup a secondary agent in Hot Standby mode. However, you can setup a second Tivoli Performance Analyzer Agent and keep it stopped, as long as the primary server is running. The secondary agent can only be started when the primary agent is stopped and disabled. The Tivoli Monitoring Administrator can perform the switch manually, or use Take Action commands. In both cases it is very important to ensure that only one agent is running at the same time. See “Configuring for Summarization and Pruning Agent high availability and disaster recovery” on page 46 for information on writing the Take Action command.

Agent deployments

When planning your installation, you need to determine how you want to deploy your agents. For very small environments, some users manually install agents on each server. For larger environments, automation tools must be used to deploy the agents and agent patches.

A key decision point is to determine which deployment software to use. The Tivoli Monitoring product has a remote deployment capability that allows you to initially deploy your operating system agents remotely, and then remotely add agents to your systems.

Each product and fix pack includes a remote deployment bundle that can be placed in the remote deployment depot for future agent distributions and patching. However, the remote deployment capability is not as efficient at distributing software as some purchasable distribution products. If you

already have an enterprise-class software distribution product like Tivoli Configuration Manager or Tivoli Provisioning Manager, you might find it more efficient to distribute the agents and patches. Tivoli Monitoring agents provide software package blocks that can be used by Tivoli Configuration Manager and Tivoli Provisioning Manager to distribute the agents.

The main advantage of using products such as Tivoli Configuration Manager and Tivoli Provisioning Manager are:

- Faster distribution times to speed large-scale deployments.
- Tivoli Configuration Manager and Tivoli Provisioning Manager can be tuned to utilize only a portion of the network bandwidth.
- Tivoli Configuration Manager and Tivoli Provisioning Manager can easily be configured for retries and tracking of success and failure.

Use the **tacmd exportBundles** command to export one or more deployment bundles to the specified export directory for use with software distribution products. For more information see the *IBM Tivoli Monitoring Command Reference*.

If you have an enterprise-class software distribution product, use it for the initial software distribution and for the deployment of larger fix packs. For interim fixes and small fix packs, the remote deployment capability might require less time to configure and utilize.

One of the most important aspects of agent deployment is the agent prerequisite preparation. Ensure that the servers are prepared with the appropriate file systems including adequate space. In addition to disk space, you must determine the user account to be used for the agent installation. By using an administrative account (administrator or root), you ease your agent deployment tasks.

If administrator or root are not allowed, then using sudo on a UNIX system is the next best choice. Without administrative authority, the installation becomes a multi-step process where the systems administrators need to be brought in to run commands such as setperm to setup the permissions. For planning purposes, plan for roughly 500 MB of disk space to allow space for the agent and historical logs.

Self-describing monitoring agents

The self-describing agent feature makes it possible for new or updated IBM Tivoli Monitoring agents to become operational after installation, without having to perform additional product support installation steps.

Self-describing agents apply version updates to other components automatically without the need to recycle your hub Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, or remote Tivoli Enterprise Monitoring Server. This feature integrates the installation of an agent with the dispersal and installation of associated product support files throughout your IBM Tivoli Monitoring infrastructure.

For more information, see “Enabling self-describing agent capability at the hub monitoring server” on page 204 and the *IBM Tivoli Monitoring Administrator's Guide*.

Prerequisite Scanner for IBM Tivoli Monitoring agents

You can perform prerequisite scanning for agents before carrying out an installation.

The two mechanisms available are a manually executed stand-alone Prerequisite Scanner and a remote prerequisite scanner facility that extends the capabilities of IBM Tivoli Monitoring's remote deployment component.

Note:

1. Only the operating system (OS) agents are enabled with prerequisite checking. If you attempt to run remote prerequisite scanning with an application agent, the operation fails because the application agent does not provide a Prerequisite Scanner.
2. On Windows platforms, the prerequisite scanner requires the Microsoft .Net Framework v2.0 to be installed.

Stand-alone prerequisite scanning

Review these topics before you run the stand-alone Prerequisite Scanner.

- “Prerequisite Scanner”
- “Running the stand-alone Prerequisite Scanner”
- “Sample Prerequisite Scanner report” on page 51

Note: Prerequisite Scanner output for both stand-alone and remote prerequisite scanning is available in English only. The Prerequisite Scanner report (results.txt file), precheck.log file, prereq_checker.bat file, and prereq_checker.sh file are also available only in English.

Prerequisite Scanner:

The Prerequisite Scanner is located on the IBM Tivoli Monitoring agent media. Platform-specific scripts are used by the prerequisite checker, so you must select the Prerequisite Scanner that is appropriate for the platform where you are executing the prerequisite scan.

The Prerequisite Scanner for **Linux** **UNIX** platforms is located in the following directory:
<ITM_Agent_Media>/unix/prereqchecker

The Prerequisite Scanner for the **Windows** platform is located in the following directory:
<ITM_Agent_Media>/WINDOWS/prereqchecker

Running the stand-alone Prerequisite Scanner: About this task

Complete the following steps to run the stand-alone prerequisite scanner:

Procedure

1. Copy the Prerequisite Scanner for the appropriate platform to the target computer. –OR–
Open a command prompt or shell environment to the Prerequisite Scanner directory on the agent image and run the Prerequisite Scanner command as described in step 3.
2. Open a command-line interface:
 - **Windows** Open a Command Prompt.
 - **Linux** **UNIX** Open a shell environment.
3. Change directory to the folder where you located the Prerequisite Scanner and run the following command:
 - **Windows** Run the prereq_checker.bat command.
 - **Linux** **UNIX** Run the prereq_checker.sh command.

The following input parameters are displayed:

```
prereq_checker.bat "<Product Code> [product version], <Product Code>
[product version]..." [PATH=<Agent install path>] [detail]
[-p <Product Code>.SECTION.NAME=VALUE pairs] [outputDir="<PRS output dir path>"]
```

Windows example: prereq_checker.bat "KNT" detail PATH=c:\ibm\itm -p SERVER=IP.PIPE://mytems:1918 outputDir="c:\PRS"

Linux **UNIX** example: ./prereq_checker.sh "KLZ" detail PATH=/opt/IBM/ITM outputDir="/tmp/prereq_check_output/" -p SERVER=mytems:1918/

Results

The following paragraphs provide a detailed description of the input parameters:

- **Product code and version pairs**

- At least one product code is required. For IBM Tivoli Monitoring, the code is the three-letter product code for the component or agent. Optionally, you can specify a version for the agent. If you do not specify a version, the Prerequisite Scanner checks the latest version available.
- You must enter the product codes in capital letters. For the Linux agent, for example, you must enter KLZ. If you enter lowercase klz, an error message is displayed.
- You can specify multiple product codes separated by commas.
- The product codes and versions must be enclosed in quotation marks, for example "KNT 06230000, KUD". In this example, the prerequisite check is made on the KNT agent version 06230000 and the latest version of the KUD agent.
- Each component or agent has a config file with the naming convention *.cfg, for example, KNT_0620000.cfg. If you specify a product code that does not have a corresponding *.cfg file, that product code is ignored.

- **[PATH=<product install path>]**

- The *PATH* parameter is optional. An example *PATH* is *PATH=D:\IBM\ITM*. If you do not specify a *PATH* parameter, the Prerequisite Scanner checks the following default IBM Tivoli Monitoring installation paths:
 - **Windows** The default path is: C:\IBM\ITM.
 - **Linux** **UNIX** The default path is: /opt/IBM/ITM.

- **[detail]**

- The *detail* parameter is optional. This flag indicates that you want to see detailed results on the screen when you run the prerequisite checker.
- Do not enclose the word *detail* in quotation marks.
- If you do not specify *detail*, then only PASS or FAIL is displayed on the screen. Omitting the parameter is a simple way to see if your system meets the prerequisites.
- See the following example output when the *detail* parameter is specified:

```
Windows OS monitoring Agent [version 06210000] :
Evaluation      PASS/FAIL      Expected Result      Result
CPU Number      PASS            1                    1
Available Memory PASS            35MB                 1.02GB
Available Disk   PASS            70MB                 1.09GB

ALL COMPONENTS :
Evaluation      PASS/FAIL      Expected Result      Result
Available Memory PASS            35MB                 1.02GB
Available Disk   PASS            70MB                 1.09GB
```

- The detailed output is written to a file called *precheck.log*, regardless of whether you specify *detail* or not. You can view this file to see the results of the Prerequisite Scanner. In addition, you can view the *result.txt* file to see the detailed steps taken by the Prerequisite Scanner.
- You can specify optional parameters by using a *-p* flag. The parameters are formatted as follows:
[-p <Product Code>.<instance>.<parameter>=<value>, <Product Code>.<instance>.<parameter>=<value>,,]

See the following example to check the ports:

```
-p SERVER=IP.PIPE://mymachine:1918
```

- **[outputDir="<fully qualified path>"]**

- You must supply a fully qualified path and place that value in quotes. The *outputDir* option is case sensitive.

Note: This script accepts the `-p` connection parameter in the same way as the `tacmd createNode` command, you can specify *SERVER*, *PROTOCOL*, *PORT*, *BACKUP*, *BSERVER*, and so on. All parameters are case sensitive, therefore you must use capital letters.

Sample Prerequisite Scanner report:

This section contains a sample Prerequisite Scanner report.

```
IBM Tivoli Prerequisite Scanner
  Version   : 1.0.17
  Build     : 2010728
  OS Name   : Microsoft Windows Server 2003, Enterprise Edition
Service Pack 2
  User Name : Administrator
Machine Info
  Machine name : DEV-TAYLOR50
  Serial Number: KLXZL5G
  OS Serial    : 69713-640-1083907-45085
```

Windows OS monitoring Agent [version 06230000]:

Property	Result	Found	Expected
=====	=====	=====	=====
# CPU	PASS	1	1
Memory	PASS	313MB	35MB
Disk	PASS	6.51GB	125MB
OS Version	PASS		

```
Microsoft Windows Server 2003, Enterprise Windows 2000 Server
Windows 2000 Advanced Server
Windows XP Professional Service Pack 1
Windows Server 2003 Datacenter Edition Service Pack 1
Windows Server 2003 Standard Edition Service Pack 1
Windows Server 2003 Enterprise Edition Service Pack 1
Windows Server 2003 R2 Enterprise x64 Edition
Windows Server 2003 R2 Standard x64 Edition
Windows Server 2003 R2 Data Center Edition x64 Edition
Windows Server 2003 Enterprise Itanium Edition
Windows Server 2008 Enterprise Edition
Windows Server 2008 Standard Edition
Windows Server 2008 Data Center Edition
Windows Server 2008 Enterprise x64 Edition
Windows Server 2008 Standard x64 Edition
Windows Server 2008 Data Center x64 Edition
```

ALL COMPONENTS :

Property	Result	Found	Expected
=====	=====	=====	=====
Memory	PASS	313MB	35MB
c:\	PASS	6667MB	125MB

Remote prerequisite scanning

Review the requirements for using remote prerequisite scanning.

You must meet the following requirements in your IBM Tivoli Monitoring infrastructure to make use of this feature:

- The command-line component known as KUE (User Interface Extensions) must be at V6.2.3 or higher. Use this component to import the agent-specific Prerequisite Scanners. The KUE component also contains the new `tacmd checkprereq` command to run the remote prerequisite scanning. You can use the `cinfo -i` command on Linux and UNIX systems or the `kincinfo -i` command on Windows systems to verify the installed level of KUE.
- The hub Tivoli Enterprise Monitoring Server must be at V6.2.3 or higher.

- For improved performance and parallelization, the `tacmd checkprereq` command uses the *SERVER* property, similar to `tacmd createnode`, to route the processing from the hub Tivoli Enterprise Monitoring Server to connect to the remote Tivoli Enterprise Monitoring Server. However, this routing occurs only if the supplied remote Tivoli Enterprise Monitoring Server is at V6.2.3 or higher. If the supplied remote Tivoli Enterprise Monitoring Server is at a version lower than V6.2.3, the remote prerequisite scanning is performed by the hub Tivoli Enterprise Monitoring Server.

Importing the agent-specific Prerequisite Scanner:

As with IBM Tivoli Monitoring agents, an agent-specific Prerequisite Scanner is delivered as a bundle in the IBM Tivoli Monitoring agent media.

The IBM Tivoli Monitoring `tacmd addbundles` command has been updated to implicitly import an agent's Prerequisite Scanner bundle into the Tivoli Enterprise Monitoring Server depot if it exists on the agent media. No additional steps must be taken beyond adding the agent bundle to the Tivoli Enterprise Monitoring Server depot, as is done for remote deployment in previous releases. The IBM Tivoli Monitoring OS agents (KNT, KUX, KLZ) are the initial agents to use the new prerequisite scanning feature.

Using remote prerequisite scanning:

IBM Tivoli Monitoring has the ability to preemptively check for required prerequisites on a remote computer. This feature uses the same functionality as the `tacmd createnode` command to transfer and run on a target computer.

By using the new **`tacmd checkprereq`** command, you can run a prerequisite scan for a specified agent on a specified target computer by including the `-h | --host` flag, or on a group of target computers by including the `-g | --deploygroup` flag.

You can also run the `tacmd checkprereq` command by using the `-n | --node` flag to perform remote prerequisite checking without endpoint credentials. Both the hub Tivoli Enterprise Monitoring Server and the Tivoli Enterprise Monitoring Agent must be at V6.2.3 or higher to use prerequisite scanning without endpoint credentials. If a prerequisite check without endpoint credentials is run on a target that has a Tivoli Enterprise Monitoring Agent older than V6.2.3, the command fails. The `-n | --node` flag allows you to run a prerequisite check for an endpoint before you attempt a deployment, without specifying endpoint credentials. You can review the prerequisite check results through the `tacmd getdeploystatus` command or the Tivoli Enterprise Portal Deploy Status workspace. If a prerequisite check fails, you can review the results log on the Tivoli Enterprise Monitoring Server in the `CANDLEHOME/logs/checkprereq_results` directory.

The following example shows the execution of the `checkprereq` command to a single computer for the Windows OS Agent. After running the command, you receive a transaction ID, which is used to track the status of the prerequisite check by using the `tacmd getdeploystatus` command or the Tivoli Enterprise Portal Deployment Status Workspace.

Example:

```
tacmd checkprereq -h mysys1.mycompany.com -t NT -p SERVER=rtems1.mycompany.com
```

Note: You are prompted for the remote login credentials. Alternatively, you can use the optional flags to supply the credentials at the command-line.

The following example shows how you can use deployment groups to perform prerequisite scanning for a group of target computers.

Example:

```

tacmd creategroup -g NewSystems -t DEPLOY

tacmd addgroupmember -g NewSystems -t DEPLOY -m
mysys1.mycompany.com -p SERVER=rtems1.mycompany.com
KDYRemote Execution and Access.Remote Execution and
AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution
and AccessPASSWORD=mypass

tacmd addgroupmember -g NewSystems -t DEPLOY -m
mysys2.mycompany.com -p SERVER=rtems2.mycompany.com
KDYRemote Execution and Access.Remote Execution and
AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution
and AccessPASSWORD=mypass

tacmd addgroupmember -g NewSystems -t DEPLOY -m
mysys3.mycompany.com -p KDYRemote Execution and
Access.Remote Execution and AccessUSERNAME=root
KDYRemote Execution and Access.Remote Execution and
AccessPASSWORD=mypassword
tacmd checkprereq -g NewSystems -t LZ

```

The steps in the previous section demonstrate the creation of a new deployment group called *NewSystems* and the addition of several deployment targets and the execution of the preemptive prerequisite scanning. The first member is using the *SERVER* property to specify that the execution should be performed by the Tivoli Enterprise Monitoring Server running on *rtems1.mycompany.com*. If this specification is not a valid Tivoli Enterprise Monitoring Server, the prerequisite scanning is performed by the hub Tivoli Enterprise Monitoring Server. The second member specifies a different executing Tivoli Enterprise Monitoring Server, and the third member does not supply the *SERVER* property, therefore the command runs on the hub Tivoli Enterprise Monitoring Server by default. When the preemptive prerequisite scanning is run, a transaction ID is provided for deployment status tracking. After performing the `tacmd checkprereq` command, and resolving any prerequisite failures so that the `tacmd checkprereq` command results in a *SUCCESS* for all target machines, you can run the `tacmd createnode` command using the same deployment group, for example `tacmd createnode -g NewSystems`.

Note: For deployment groups with mixed platform types such as a combination of Linux and UNIX systems, and Windows systems, you can specify the product as *ITM*. In the case of the OS agents, the correct OS agent type is determined based on the platform type discovered when the connection is established. Even when the incorrect product code is supplied for the OS agents, the correct platform type is detected. This is a special case that applies to OS agents.

After running the preemptive prerequisite scanning command, the example run can result as follows. In the following example, you can see that two of the endpoints successfully passed the prerequisite scan, however the *mysys2.mycompany.com* endpoint did not pass.

```

C:\IBM\ITM\bin>tacmd getdeploystatus -g 12829182384530000000000041

Transaction ID : 12829182384530000000000041
Command       : CHECKPREREQ
Status        : SUCCESS
Retries       : 0
TEMS Name     : RTEMS1
Target Hostname: mysys1.mycompany.com
Platform      : li6263
Product       : LZ
Version       : 062300000
Error Message  : KDY4001I: The prerequisite checking operation was a success.

Transaction ID : 12829182384530000000000041
Command       : CHECKPREREQ
Status        : FAILED
Retries       : 3
TEMS Name     : RTEMS2
Target Hostname: mysys2.mycompany.com

```

```

Platform      : li6263
Product       : LZ
Version       : 062300000
Error Message : KDY4003E: A failure occurred while checking for
required prerequisites for bundle LZ on the
target host mysys2.mycompany.com.
Review the results from the prerequisite checking execution and
resolve the issues found on the endpoint.

Transaction ID : 1282918238453000000000041
Command       : CHECKPREREQ
Status        : SUCCESS
Retries       : 0
TEMS Name     : HUB_TEMS
Target Hostname: mysys3.mycompany.com
Platform      : li6263
Product       : LZ
Version       : 062300000
Error Message : KDY4001I: The prerequisite checking operation was a success.

```

To correct the failures found on the mysys2.mycompany.com endpoint you can review the results log that was sent back to the executing Tivoli Enterprise Monitoring Server (RTEMS2 in this case). The results file is located in the CANDLEHOME/logs/checkprereq_results directory. The naming format of the results file is <SUCCESS|FAILED>_<host>_<transaction ID>.txt. In this instance, the name of the results file is FAILED_mysys2.mycompany.com_1282918238453000000000041.txt. The preceding example shows the output of the prerequisite scan report. To review the results of both *SUCCESS* and *FAILED* prerequisite checks, you must specify the -c | -- collectall flag when running the tacmd checkprereq command. By default, all *FAILED* prerequisite check results are sent back to the executing Tivoli Enterprise Monitoring Server.

CLI and remote deployment prerequisite scanning:

You can use the remote deployment prerequisite scanning feature as part of a tacmd addSystem, tacmd createnode, or tacmd updateAgent command execution when deploying an agent.

Use the command in this way to run a prerequisite check in a single command-line execution. You have the option of allowing the deployment to continue only if the check is successful, or to ignore the results and attempt the deployment even if the Prerequisite Scanner fails. The valid options are: COLLECTALL, EXECPREREQCHECK, and IGNOREPREREQCHECK. The values are to be specified in KEY=VALUE format.

Note: The prerequisite scanning feature is available only on Tivoli Enterprise Monitoring Agents at Tivoli Monitoring V6.2.3 or later. If prerequisite scanning is enabled during a tacmd addSystem or tacmd updateAgent command execution for a target that has a Tivoli Enterprise Monitoring Agent older than V6.2.3, the command fails with an error message.

An agent-specific Prerequisite Scanner is delivered as a bundle in the IBM Tivoli Monitoring agent media. The IBM Tivoli Monitoring tacmd addbundles command is used to import an agent's Prerequisite Scanner bundle into the Tivoli Enterprise Monitoring Server depot if it exists on the agent media. No additional steps must be taken beyond adding the agent bundle to the Tivoli Enterprise Monitoring Server depot as was done for remote deployment in previous releases. Next, issue a tacmd addSystem, tacmd createnode, or tacmd updateAgent command with prerequisite scanning enabled (EXECPREREQCHECK=Y). The following outcomes are possible:

- If the tacmd addSystem command runs successfully, a new application agent is installed on the endpoint and connected back. If the tacmd createnode command runs successfully, a new OS agent is installed on the endpoint and connected back. If the tacmd updateAgent command runs successfully, a new version of the agent is installed on the endpoint and connected back.
- If the prerequisite check fails, review the prerequisite check results log located on the deploying Tivoli Enterprise Monitoring Server in the CANDLEHOME/logs/checkprereq_results directory.

- The prerequisite check might fail but the installation still continues (due to the IGNOREPREREQCHECK option) and the installation fails. In this scenario, the data is received through the tacmd getdeploystatus command or the Tivoli Enterprise Portal Deploy Status workspace.

Background information about agent autonomy

By default, agents are configured for autonomous operation (in other words, the default value for configuration parameter IRA_AUTONOMOUS_MODE is Y). This means that when the connection to the Tivoli Enterprise Monitoring Server is lost or the monitoring server otherwise becomes unavailable, the agent continues to run all situations that can be evaluated exclusively at the agent.

Situations that become true while disconnected from the monitoring server are stored persistently and are uploaded to the monitoring server when reconnected. If the agent was starting up when the communications break occurred, startup processing continues.

If you do not want autonomous agent operation, you can disable it by setting environment variable IRA_AUTONOMOUS_MODE to N in the agent configuration. In this case you must also ensure variable CT_CMSLIST is both specified and not blank; otherwise the agent will not start.

To configure an agent for fully autonomous operation (in other words, it runs without a connection to a monitoring server), the following agent configuration is required:

1. The CT_CMSLIST variable must not be set (that is, it should have no value).
2. At least one protocol must be defined in the KDC_FAMILIES variable.
3. The IRA_AUTONOMOUS_MODE variable does not need to be set, as its default value is Y. If IRA_AUTONOMOUS_MODE is currently set to N, it must be either changed to Y or removed completely.

Note:

1. If you respond NO to question Will this agent connect to a TEMS? when running the UNIX installer, these parameters get set correctly for you.
2. In all cases, including fully autonomous mode, at least one active protocol *must* be defined by using the KDC_FAMILIES environment variable. If no protocols are defined, the agent will not start.

Event forwarding from autonomous agents

By incorporating the Event Integration Facility (EIF) into the autonomous agents, they can forward events directly to Netcool/OMNIbus via the Secure Sockets Layer (SSL) protocol. Available for the Windows, Linux, UNIX, and z/OS platforms, this integration provides another event emitter (similar to the SNMP trap emitter) that can emit EIF-format events directly from an autonomous agent to the event-management facility your site uses.

Currently, EIF event notification is supported only for locally defined private situations—those defined via a private configuration file local to the monitoring agent. (Events for enterprise situations must still be forwarded from the hub Tivoli Enterprise Monitoring Server.) These EIF events provide similar function and format to those provided by the event forwarder on the hub monitoring server, including the ability to send situation status. There is no support for dynamic refresh of changed EIF events (that is, event maps and event destinations); for such changes to be effective, you must recycle the agent.

You can customize agent-generated EIF events via a product-provided or user-defined event mapping file, which specifies these elements:

- The attributes included in the event.
- Custom message slot text with variable substitution.
- Mapping by situation name.
- Mapping by attribute group used in the situation.

Events generated will have a source slot value of ITM Agent:Private Situation to distinguish them from events originated at the monitoring server, which have a source slot value of ITM.

You can enable or disable event forwarding by setting IRA_EVENT_EXPORT{EIF to either Y or N in the event mapping file. Event forwarding is disabled automatically when no valid EIF event destination is defined. Note that a single event can be sent to up to five different destinations; these can be a mixture of both TEC and non-TEC event receivers.

A heartbeat function lets you use the event-forwarding feature to send events to either Tivoli Enterprise Console or Netcool/OMNIBus that notify you if the agent is online and running. The heartbeat interval determines how often the agent generates this event; it is configurable via the event mapping file. These EIF events, whose classname is ITM_Heartbeat, contain a slot called interval whose value is the heartbeat interval. You can customize the IBM-provided heartbeat rules or write your own to handle heartbeat events as your site's needs dictate.

For complete information about defining your own event mapping file, see the *IBM Tivoli Monitoring Administrator's Guide*.

Note: Existing TEC and OMNIBus event-synchronization rules function as usual, but bidirectional interactions with IBM Tivoli Monitoring are not possible, as the events do not originate with the Tivoli Enterprise Monitoring Server. However, the TEC and OMNIBus event receivers will still be able to update event status for events sent from autonomous agents.

Agentless monitoring versus monitoring agents

This section can help you to understand the characteristics of operating system agents and agentless monitors, as well as the differences between them.

IBM Tivoli Monitoring provides operating system (OS) agents that monitor the availability and performance of the computers in your monitoring environment. An example of an OS agent is the monitoring agent for Windows, which can monitor Windows XP, Windows 2003, and Windows 2008 operating systems. These full-function OS agents must reside on the same computers they are monitoring.

IBM Tivoli Monitoring also provides *agentless monitors*. An agentless monitor is a standard Tivoli Monitoring agent that can monitor the operating system running on multiple remote nodes that do not have the full-function OS agents running on them. An agentless monitor obtains data from nodes it is monitoring via a remote application programming interface, or API—in this case, SNMP, CIM, or WMI—running on the node being monitored. Since these interfaces provide information about either operating system functions or base application functions, no IBM Tivoli Monitoring component need be installed or deployed on the monitored node.

API Function

SNMP

The Simple Network Management Protocol is a TCP/IP transport protocol for exchanging network management data and controlling the monitoring and operation of network nodes in a TCP/IP environment.

CIM The Common Information Model is an XML-based standard for defining device and application characteristics so system administrators and management programs can monitor and control them using the same set of tools, regardless of their differing architectures. CIM provides a more comprehensive toolkit for such management functions than the Simple Network Management Protocol.

WMI Microsoft's Windows Management Instrumentation API provides a toolkit for managing devices and applications in a network of Windows-based computers. WMI provides data about the status of local or remote computer systems as well as the tools for controlling them. WMI is included with the Windows XP and Windows Server 2003 and 2008 operating systems.

These APIs are supported by the Agent Builder, which enables you to build custom agentless monitoring solutions that are separate from the agentless monitors available on the Tivoli Monitoring installation media and that provide additional function.

Since an agentless monitor is a standard Tivoli Monitoring agent, it collects data and distributes it to a Tivoli Enterprise Monitoring Server and then on to a Tivoli Enterprise Portal Server. It also takes advantage of the various features of the IBM Tivoli Monitoring product, such as Tivoli Enterprise Portal workspace views, situations, remote deployment of the agentless monitors, policies, and so on. Detailed information can be found in the user's guide for each agentless monitor; see Table 7 on page 61.

Agentless monitoring does not provide the kind of deep-dive information your site may need for its core business servers; however, it does allow a small set of centralized servers to supervise the health of the operating nodes in your environment. There are five types of agentless monitors that cover the Windows, AIX, Linux, HP-UX, and Solaris environments.

The agentless monitors are multi-instance agents. After installing or deploying an agentless monitor on a machine, additional instances can be created via configuration. Each instance can communicate with up to 100 remote nodes.

Each type of agentless monitor can run on additional platforms beyond the type of platform it monitors. For example, the agentless monitor for Windows (which monitors only Windows operating systems) can run on any of the supported platforms: Windows, AIX, Solaris, HP-UX, Linux.

Specific operating system releases that a particular agentless monitor can monitor are detailed in Table 6 on page 59. Check the user's guide for each agentless monitor regarding platform-specific requirements for the operating systems that agentless monitors can run with.

A computer that has one or more agentless monitors running on it is referred to as an agentless monitoring server. Each server node can support up to 10 active agentless monitor instances, in any combination of agentless monitor types; for example, 2 AIX, 2 HP-UX, 2 Linux, 2 Solaris, 2 Windows; or 4 Windows, 3 AIX, 3 Linux; or 5 Windows, 5 Solaris; or 10 HP-UX. Each instance can communicate with up to 100 remote nodes, which means a single agentless monitoring server can support as many as 1000 monitored systems (10 instances * 100 remote nodes per instance). By adding more server nodes, the number of monitored nodes increases into the thousands.

Figure 8 on page 58 illustrates the architecture of an IBM Tivoli Monitoring environment that employs agentless monitoring.




Figure 8. Architecture of agentless monitoring

Agentless technology provides lightweight OS monitoring that targets key metrics along with basic situations meant to satisfy simple monitoring needs. Agentless monitoring provides speedy implementation and minimum agent deployment, including the deployment of updates; however, the need to poll the monitored node to retrieve its monitoring data increases network traffic, and real-time data availability is impacted both by the network delay and the reliance on polling. In addition, the implementation of Take Action commands for command and control is more powerful with the full-function agents than for agentless technology.

Key operating system metrics returned:

- Logical and physical disk utilization.
- Network utilization
- Virtual and physical memory
- System-level information
- Aggregate processor utilization
- Process availability

Default situations are provided for:

- Disk utilization
- Memory utilization
- CPU utilization
- Network utilization

You can use these situations as is or as models for custom situations that meet your site's specific needs.

The agentless monitors monitor the distributed operating systems listed in Table 6 on page 59. You can configure different data collectors for these environments, as shown.

Table 6. Data collectors usable with the various agentless monitors and releases supported

Agentless monitor	Product code	Data collectors supported	Operating system releases monitored
Agentless Monitoring for Windows OS	R2	<ul style="list-style-type: none"> • WMI¹ • Performance Monitor (PerfMon)¹ • Windows event log¹ • SNMP V1, V2c, V3 	<p>Monitors the following Windows releases:</p> <ul style="list-style-type: none"> • Windows Server 2003 Standard Edition (32 bit) with SP1 or higher • Windows Server 2003 Standard Edition (32 bit) with R2 SP1 or higher • Windows Server 2003 Enterprise Edition (32 bit) with SP1 or higher • Windows Server 2003 Enterprise Edition (32 bit) with R2 SP1 or higher • Windows Server 2003 Datacenter Edition (32 bit) with SP1 or higher • Windows Server 2003 Datacenter Edition (32 bit) with R2 SP1 or higher • Windows 2003 Standard Edition (64 bit) with R2 SP2 or higher • Windows 2003 Enterprise Edition (64 bit) with R2 SP2 or higher • Windows Server 2003 Datacenter Edition (64 bit) with R2 SP2 or higher • Windows 2003 Server Enterprise Edition on Itanium2 (IA64) with R2 SP2 or higher • Windows Server 2008 Standard Edition (32 bit) • Windows Server 2008 Datacenter Edition (32 bit) • Windows Server 2008 Enterprise Edition (32 bit) • Windows Server 2008 Standard Edition (64 bit) • Windows Server 2008 Enterprise Edition (64 bit) • Windows 2008 Enterprise Edition on Itanium2 (IA64) • Windows Vista Enterprise, Business, and Ultimate (32 bit) • Windows Vista Enterprise, Business, and Ultimate (64 bit) <p>Note: IA64 machines running Windows are not supported.</p>
Agentless Monitoring for AIX OS	R3	SNMP V1, V2c, V3	<p>Monitors the following AIX releases:</p> <ul style="list-style-type: none"> • AIX V6.x (64 bit)

Table 6. Data collectors usable with the various agentless monitors and releases supported (continued)

Agentless monitor	Product code	Data collectors supported	Operating system releases monitored
Agentless Monitoring for Linux OS	R4	SNMP V1, V2c, V3	Monitors the following Linux releases running on xSeries®, pSeries, and zSeries® machines: <ul style="list-style-type: none"> • RedHat Enterprise Linux 5 Intel (32 bit) • RedHat Enterprise Linux 5 on x86-64 • RedHat Enterprise Linux 5 on Itanium 64 bit • RedHat Enterprise Linux 5 on iSeries® and pSeries • RedHat Enterprise Linux 5 on zSeries (64 bit) • SuSE Linux Enterprise Server 9 Intel (32 bit) with SP3 or later • SuSE Linux Enterprise Server 9 on x86-64 (64 bit) with SP3 or later • SuSE Linux Enterprise Server 9 on Itanium (64 bit) with SP3 or later • SuSE Linux Enterprise Server 9 for iSeries and pSeries with SP3 or later • SuSE Linux Enterprise Server 9 for zSeries (31 bit) with SP3 or later • SuSE Linux Enterprise Server 9 for zSeries (64 bit) with SP3 or later • SuSE Linux Enterprise Server 10 Intel (32 bit) • SuSE Linux Enterprise Server 10 on x86-64 (64 bit) • SuSE Linux Enterprise Server 10 on Itanium (64 bit) • SuSE Linux Enterprise Server 10 for iSeries and pSeries (64 bit) • SuSE Linux Enterprise Server 10 for zSeries (64 bit)
Agentless Monitoring for HP-UX OS	R5	SNMP V1, V2c, V3	Monitors the following HP-UX releases: <ul style="list-style-type: none"> • HP-UX 11i v1 (B.11.11) (32/64) on PA-RISC • HP-UX 11i v2 (B.11.23) (64 bit) on PA-RISC • HP-UX 11i v3 (B.11.31) (64 bit) on PA-RISC • HP-UX 11i v2 (B.11.23) on Integrity (IA64) • HP-UX 11i v3 (B.11.31) on Integrity (IA64)
Agentless Monitoring for Solaris OS	R6	<ul style="list-style-type: none"> • CIM-XML • SNMP V1, V2c, V3 	Monitors the following Solaris releases: <ul style="list-style-type: none"> • Solaris V8 (SPARC) (32/64bit) • Solaris V9 (SPARC) (32/64bit) • Solaris V10 (SPARC) (32/64 bit) • Solaris V10 (x86-64) (64 bit) • Solaris V10 (Opteron) (64 bit)
¹ To use one of the native Windows data collectors (WMI, PerfMon, the event log), the agentless monitoring server must run under Windows.			

IBM recommends that you deploy a full-feature operating system agent to each agentless monitoring server to watch the CPU, memory, and network consumption of the agentless monitors themselves.

Deployment options for agentless monitors

As with other IBM Tivoli Monitoring agents, you can either install agentless monitors or deploy them from your site's deployment depot.

When installing Tivoli Monitoring, you can add the agentless monitors to your site's deployment depot, as shown in Figure 9 on page 61.

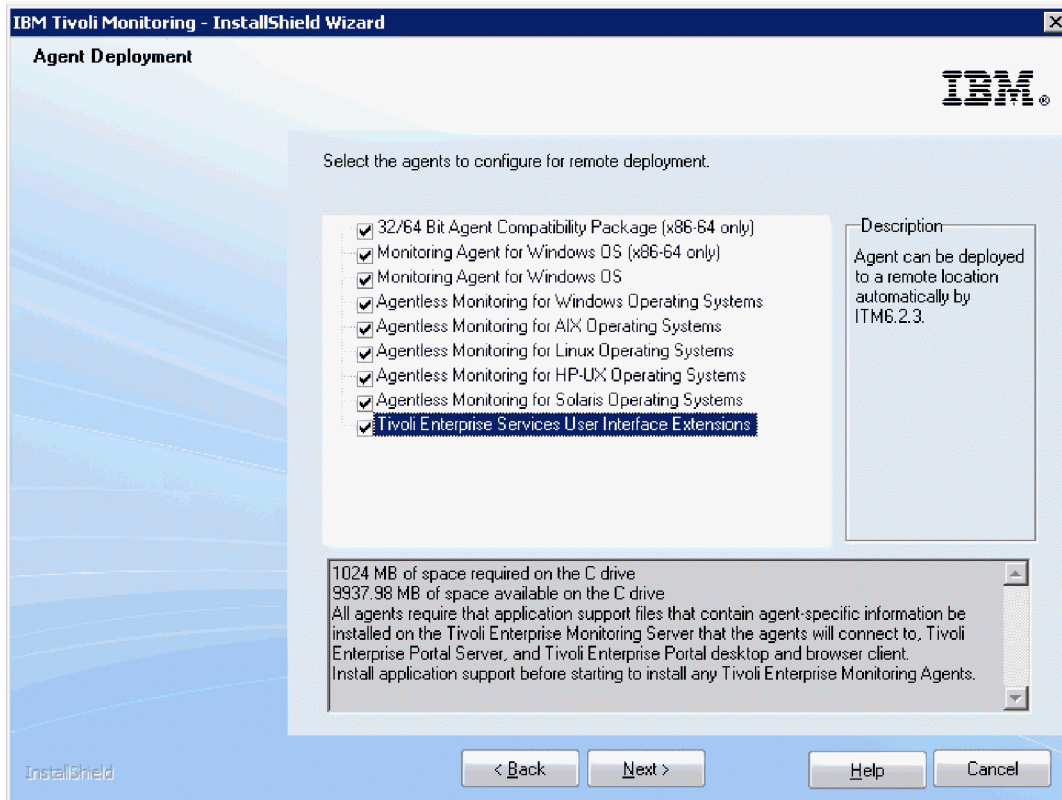


Figure 9. Adding agentless monitors to the deployment depot

You also have the full range of remote-deployment options, as explained in Chapter 4, “Deployment phase,” on page 97, at your disposal when planning how best to deploy agentless monitors across your environment. These include:

- The Tivoli Enterprise Portal's deployment features.
- The tacmd CLI commands.

As required for the deployment of any IBM Tivoli Monitoring agent, remote deployment of an agentless monitor to an agentless monitoring server requires that an OS agent be running on that machine. For example, if the agentless monitor runs on an AIX operating system, the IBM Tivoli Monitoring AIX agent must first be running on it to remotely deploy that agentless monitor. In addition, the OS agent is required to configure a server's agentless monitors via the Tivoli Enterprise Portal.

The agentless monitors are included on the same Agents DVD as the traditional OS agents.

Documentation resources for agentless monitoring

This section lists the agentless monitors user's guides.

Table 7 lists the IBM Tivoli Monitoring manuals that detail the configuration and usage of the agentless monitors.

Table 7. User's guides for the agentless monitors

Title	Document number
IBM Tivoli Monitoring Agentless Monitoring for Windows Operating Systems User's Guide	SC23-9765
IBM Tivoli Monitoring Agentless Monitoring for AIX Operating Systems User's Guide	SC23-9761
IBM Tivoli Monitoring Agentless Monitoring for Linux Operating Systems User's Guide	SC23-9762
IBM Tivoli Monitoring Agentless Monitoring for HP-UX Operating Systems User's Guide	SC23-9763

Table 7. User's guides for the agentless monitors (continued)

Title	Document number
IBM Tivoli Monitoring Agentless Monitoring for Solaris Operating Systems User's Guide	SC23-9764

Problem-diagnosis tools available for agentless monitoring

Log files are available on both the remote system being monitored and the agentless monitoring server that is doing the polling.

The following agent log files are available in these locations on the monitoring server:

- **Windows** C:\IBM\ITM\TMAITM6\logs
- **Linux** **UNIX** /opt/IBM/ITM/logs

On the remote system, the following log resources are at your disposal, depending on the type of system being monitored and the monitoring API used:

- **Windows** event logs
- **Linux** **UNIX** SNMPD or CIM log

Mainframe users

Mainframe environments have some unique considerations. There are features that are available only when running a z/OS hub monitoring server and features that are available only when running a distributed hub monitoring server.

This section outlines those considerations so that you can make the best choice for your environment.

Unique z/OS hub monitoring server features

The z/OS hub monitoring server allows you to take advantage of RACF® authentication. However, the OMEGAMON for MQ Configuration product has some specific integration with RACF that requires a z/OS hub in order to take advantage of RACF authentication within the OMEGAMON for MQ Configuration product.

The z/OS hub does not provide the Hot Standby feature. High availability is achieved using a movable hub solution as described in *Configuring the Tivoli Enterprise Monitoring Server on z/OS* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm).

Unique distributed hub monitoring server features

Three features that are provided on a distributed hub monitoring server environment are not available on z/OS hub monitoring server environments.

- Remote deployment
- Tivoli Enterprise Monitoring Automation Server and its Performance Monitoring service provider
- Hot Standby

The distributed hub monitoring server has a feature called Hot Standby to assist with high availability and disaster recovery scenarios. Many users choose not to use the Hot Standby feature and instead deploy OS Clusters for high availability and disaster recovery.

Linux on z/VM systems

Many mainframe users run Linux on their z/VM systems. Many different Tivoli Monitoring components can be installed in Linux on z/VM environments, including the monitoring server, portal server, monitoring agent and warehouse-related components. Each Linux environment can be configured with monitoring server or portal server software, or both.

Multi-hub environments

Large users who go beyond the limits of a single hub monitoring server environment must consider additional factors. Tivoli Monitoring V6.3 has been tested with environments of up to 20000 agents. Some users need multiple hub monitoring servers to handle the tens of thousands of agents in their environment.

Following are some considerations to take into account when deploying multiple hub monitoring servers.

Sharing a warehouse database

You can share a single warehouse database with multiple hub monitoring servers, but there are additional considerations when choosing this deployment option. First, you must take into account scalability of the Warehouse Proxy and Summarization and Pruning Agents. Use the Warehouse load projection spreadsheet, which can be found in the IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>) by searching for “warehouse load projections” or the navigation code “1TW10TM1Y”.

With multiple hubs and up to 20000 agents, you increase the likelihood of exceeding the capacity of the Warehouse. Be aware of how much data you are collecting to ensure that you do not exceed the capacity of the warehouse database. To get an idea of the capacity of the Summarization and Pruning agent with your warehouse database, consider using the measurement approach discussed in “Locating and sizing the Summarization and Pruning Agent” on page 39.

In addition to scalability, there are specific deployment requirements when a warehouse database is shared between hub monitoring servers. First, you can run only one Summarization and Pruning Agent in only one of the two monitoring server environments. The single Summarization and Pruning Agent is responsible for summarizing and pruning the data for all of the data in the Warehouse. The summarization and pruning configuration settings are maintained by the portal server that is specified in the Summarization and Pruning Agent configuration dialog.

Due to the complexity and potential scalability issues of sharing a warehouse database across multiple hub monitoring servers, you might want to maintain multiple warehouse databases. To build reports across the databases, use Federation capabilities or create a data mart that merges that content from multiple warehouse databases.

You cannot set up different summarization and pruning schedules for each of the hub monitoring server environments. In addition, you must also ensure that the hub with the Summarization and Pruning Agent is patched and maintained so that it is a superset of the two monitoring servers. If you install the database agents in one hub, then you must install the application support for the database agents on the hub monitoring server and portal server in the hub environment with the Summarization and Pruning Agent. If you install a fix pack on one hub, then you must ensure that it is also installed on the hub with the Summarization and Pruning Agent, which ensures that the Summarization and Pruning Agent is aware of all attribute groups and attributes that can be collected.

Sharing customization

When using multiple hubs, most customization can be shared between the two hub environments. Customization includes situations, workflow policies, workspaces, and managed systems lists. There are tacmd CLIs commands that can be used to perform bulk imports and exports of situations, workflow policies, and workspaces. For details on the tacmd CLIs, see the *IBM Tivoli Monitoring Command Reference*. Most of the customization can be cleanly exported from one monitoring server environment to another monitoring server environment using tools that are identified in “Maintaining an efficient monitoring environment” on page 121.

Dashboard environments and Authorization Policy Server

In the V6.3 release of IBM Tivoli Monitoring, you must set up a separate dashboard environment for each Hub monitoring server (unless it is a standby Hub for the Hot Standby feature). A dashboard environment consists of the Dashboard Application Services Hub and monitoring applications such as IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, or custom monitoring dashboards.

A dashboard environment can also include the Authorization Policy Server and the `tivcmd` CLI. If all of your authorization policies are going to be different for each dashboard environment, or each environment is going to have a different set of policy administrators: you should install the Authorization Policy Server with the Dashboard Application Services Hub in each of your dashboard environments. However, if some of your authorization policies are common across monitoring domains, you can use a single Authorization Policy Server instance to manage the policies for multiple domains. (A domain consists of the Hub monitoring server, its standby Hub monitoring server, remote monitoring servers, the portal server, a Dashboard Application Services Hub, and monitoring agents.) For this scenario, install the Authorization Policy Server on the Dashboard Application Services Hub for one of your domains and use the `tivcmd` CLI to create authorization policies for all of your domains using this single instance of the policy server. (As an alternative, you can install the Authorization Policy Server with a Dashboard Application Services Hub that is used solely for the purpose of creating authorization policies.) Configure the portal servers for all of your domains to retrieve authorization policies from the single instance of the Authorization Policy Server when you enable authorization policies. All of your Dashboard Application Services Hubs and portal servers must be configured to use the same LDAP user repository so that the Authorization Policy Server can validate that roles are being assigned to users and user groups in this common repository.

If you have some authorization policies that should be domain specific, you can specify the domain name when using the `tivcmd` CLI to create a permission. By default, the provider ID of the dashboard data provider is used as the domain name and its format is `itm.<Hub monitoring server name>`. However, you can customize a domain override value when you configure the portal server. If you do not specify a domain name when you create a permission then the permission applies to all domains that the authorization policies are distributed to.

Figure 10 on page 65 shows how to setup multiple domains that share a common set of authorization policies.

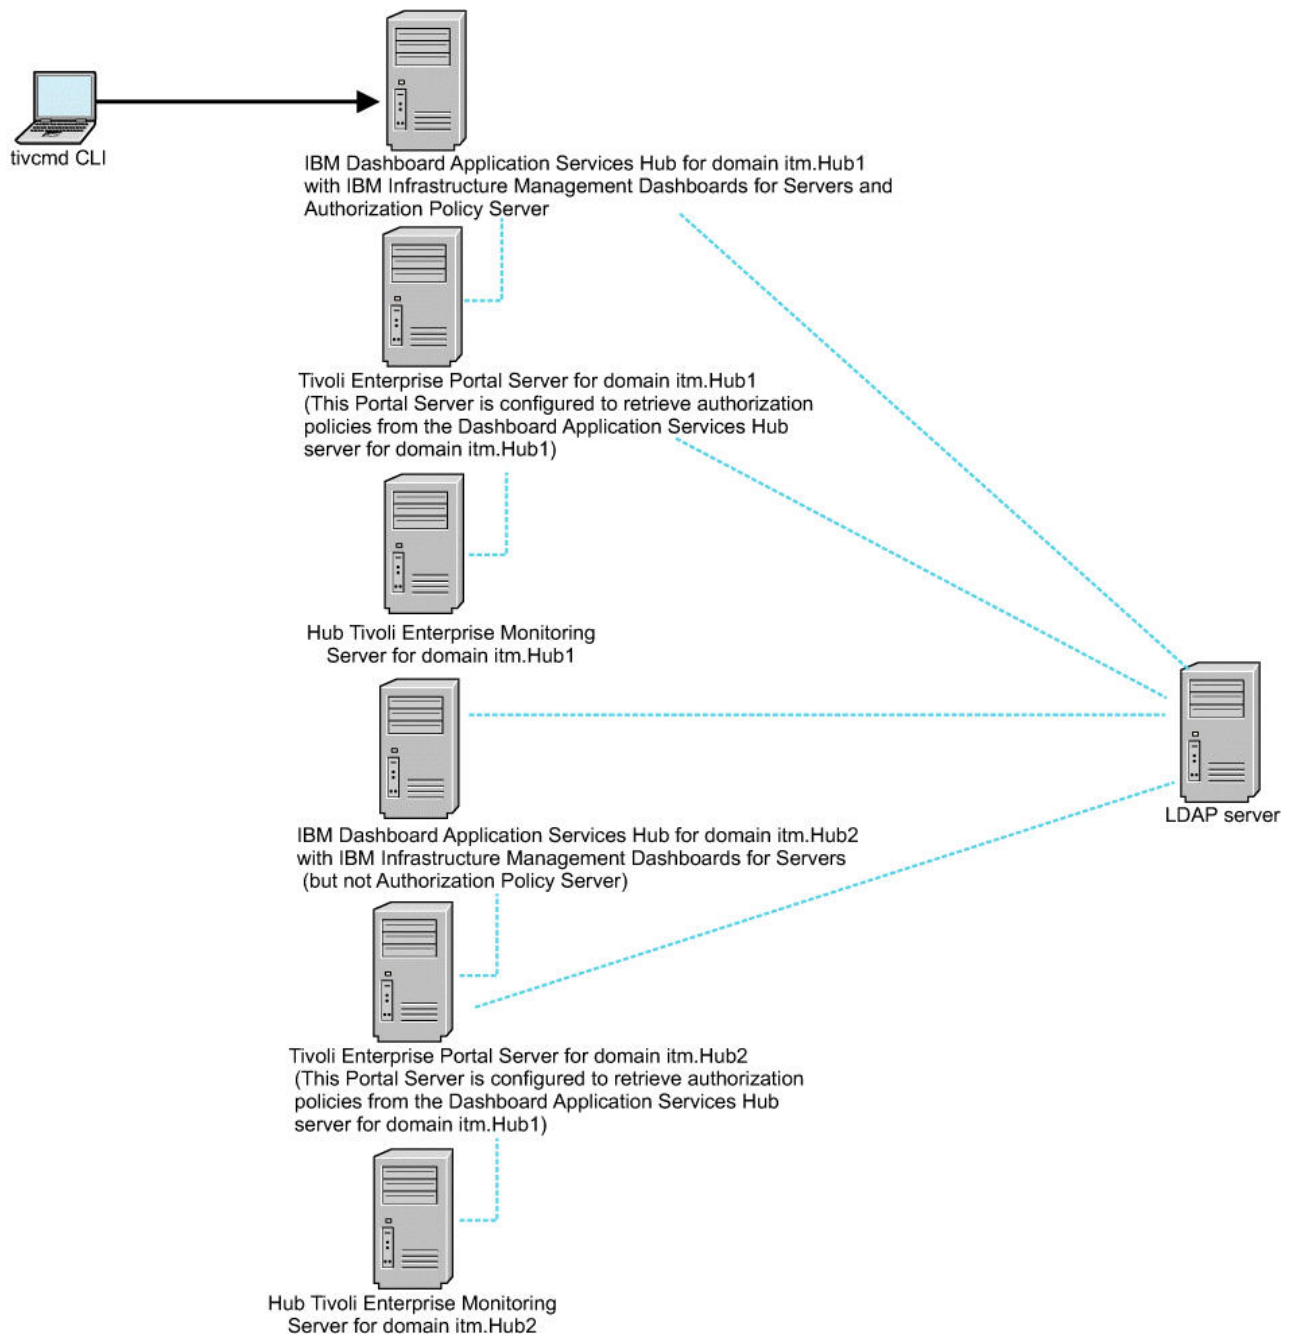


Figure 10. Multiple domains with a common set of authorization policies

See the *IBM Tivoli Monitoring Administrator's Guide* chapter for authorization policies and the *IBM Tivoli Monitoring Command Reference* for more information on creating domain specific authorization policies.

Accelerating your custom monitoring

There are ways of accelerating the creation and deployment of custom monitoring solutions.

First, you can access many solutions already available in the IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>). Check the Tivoli Integrated Service Management Library site before creating any custom solutions.

Planning and project management

The Project Management and Planning phase should be led by a certified project manager working with a deployment architect. The project manager ensures that all required tasks are identified and documented. Thus, providing a framework for tracking the progress of the deployment.

Before you begin the implementation phase of your deployment, you must develop a detailed project plan. Chapter 4, “Deployment phase,” on page 97 helps you understand the detailed steps required to successfully install and configure your Tivoli Monitoring environment. This section describes some additional tasks that must take place before the initial implementation takes place.

The most important task to complete before proceeding to the implementation phase is to identify and document the key requirements. These requirements are typically business objectives that have been identified by senior management. After the business objectives have been identified, the Tivoli Monitoring requirements and implementation details can be created by architect. These details include the following:

- High-level design for the Tivoli Monitoring components.
- Naming conventions for situation, workspaces, queries, managed systems lists, and Take Actions commands.
- Any changes to existing production environments such as firewall changes, and server prerequisites.
- Monitoring strategy

The monitoring strategy should specify exactly what will be monitored, and how; what situations and automation policies will be created; what events will be forwarded to other event management systems; and similar details. Working out the specifics of the monitoring strategy is perhaps the most challenging aspect of planning.

The details provided in this guide are intended to help the architect design the right solutions to meet the needs of all Tivoli Monitoring stake holders, given the constraints of the environment and the business objectives.

Estimating deployment tasks

The following estimating approach provides details of each logical task so that the estimates can be applied to any piece of work.

No distinction is made between the tasks carried out in test and the tasks carried out in production in terms of how long the tasks take to complete. The only difference is that in production there are more barriers to starting the work because of the need for access to user IDs, firewalls, and change control.

In each section a description of the tasks necessary to deploy those components is provided.

- “Install server components on Windows and UNIX” on page 67
- “Install server components on z/OS” on page 67
- “Install data warehousing components” on page 67
- “Install and configure event integration components” on page 68
- “Install and configure monitoring agents” on page 68
 - “Self-describing monitoring agents” on page 69
- “Setting up LDAP and single sign-on for your dashboard environment” on page 69
- “Creating authorization policies” on page 70
- “Setting up situation-based monitoring” on page 70
- “Creating workflow policies” on page 70
- “Creating workspaces” on page 70
- “Creating and deploying Agent Builder Agents” on page 70

- “Transferring skills” on page 71
- “Scheduling the initial deployment” on page 71
- “Scheduling for fix packs” on page 71

Install server components on Windows and UNIX

The first step is to install the Tivoli Monitoring server components. In general, for planning purposes, allow one day for this task.

In some cases, installing the server components takes longer than one day due to change control processes or lack of access to key personnel such as database administrators and network administrators. By arranging for key personnel to be available before beginning the installation, you ensure a timely deployment.

Tivoli Monitoring includes the following server components for a typical monitoring environment:

- DB2
- Monitoring server
- Portal server

Note: Although IBM Java is included with the Tivoli Monitoring distribution, special considerations might need to be made if there are existing versions of Java required on the server where Tivoli Monitoring is to be installed. Thus, Java might be a specific component that also needs to be addressed as part of your planning process.

The following optional server components can also be installed:

- Tivoli Enterprise Monitoring Automation Server: This component is installed on the same system as the Hub monitoring server. If you plan to use this component, you also must install these Jazz for Service Management components: Registry Services (required) and Security Services (optional).
- Dashboard Application Services Hub: This is a Jazz for Service Management component that provides the platform for installing monitoring dashboard applications such as IBM Infrastructure Management Dashboards for Servers and the Authorization Policy Server.

If you are installing these optional components, allow one additional day to get them installed.

Install server components on z/OS

Most users with a z/OS monitoring server deploy both z/OS agents and a monitoring server. The estimates in this section include the time requirements for installing both the agents and the monitoring server components.

In general, you need to allow a day for loading the software off tape and performing the receive, apply, and accept processing. The elapsed time depends on the number of products and the amount of processing power allocated to the installation LPAR, typically a test environment. If you have more than two products, then allow more time. As a rule of thumb allow half a day per product with a minimum of one day.

For more information search for “SMP/E installation on z/OS” or navigation code “1TW10TM3M” at the IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>).

Install data warehousing components

In general, allow a day to install and configure the Warehouse Proxy Agent and the Summarization and Pruning Agent.

The installation is straightforward but then you need to verify that data is being collected and summarized successfully, which has inherent delays because the minimum warehouse interval is one hour and there is no way to force the warehousing of data.

If you want to run reports against historical data, you can install IBM Dashboard Application Services Hub with Tivoli Common Reporting and install monitoring agent reporting packages. If you are planning to use IBM Dashboard Application Services Hub for monitoring dashboards, you can install Tivoli Common Reporting into this same instance. You should allow one day to get the reporting environment setup after you have verified your data warehouse environment.

Install and configure event integration components

Allow one day to install the event synchronization component and configure event forwarding to either Tivoli Enterprise Console or Netcool/OMNIBus. This estimate includes time to confirm that events are being forwarded to Tivoli Enterprise Console or Netcool/OMNIBus and synchronized back to Tivoli Monitoring.

Because this step requires the installation of new baroc files for events with Tivoli Monitoring V6.3 (Table 8 lists the update history for these files), you must plan for at least one recycle of your Tivoli Enterprise Console server. Since the installation of new baroc files for Tivoli Monitoring requires that the Tivoli Enterprise Console Server be recycled, users must plan for an outage in their Tivoli Enterprise Console Server. This means you must use your normal change control processes to schedule a change window. The estimated one-day duration to implement the Tivoli Enterprise Console changes does not include creating any custom rules or baroc files. The one day includes implementing the default baroc files included with Tivoli Monitoring.

When using the Tivoli Enterprise Console, you need a baroc file for each type of agent. For the packaged agents, a baroc file is automatically installed on your monitoring server. The baroc files are placed in the `CANDLE_HOME\cms\TECLIB` directory.

Table 8. Update history for the baroc files for IBM Tivoli Monitoring agents and components

IBM Tivoli Monitoring agent/component	Last updated for which release?
Tivoli Enterprise Monitoring Server (kib.baroc)	V6.3
Remote deployment (kdy.baroc)	V6.2.1
Warehouse Proxy Agent (khd.baroc)	V6.3
Summarization and Pruning Agent (ksy.baroc)	V6.3
Windows agent (knt.baroc)	V6.3
Linux agent (klz.baroc)	V6.3
UNIX agent (kux.baroc)	V6.3
Unix Log Alert agent (kul.baroc)	V6.2.1
IBM i agent (ka4.baroc)	V6.3
Tivoli Performance Analyzer (kpa.baroc)	V6.2.3 fix pack 1
Tivoli Performance Analyzer OS agent domain (kp3.baroc)	V6.2.3 fix pack 1

Note: File `omegamon.baroc` contains the base event class definitions for all Tivoli Monitoring events; it is automatically installed on Tivoli Enterprise Console when event synchronization is installed.

Install and configure monitoring agents

The time required to install and configure an agent differs depending on the agent. Using operating system agents as a baseline, you should be able to install a monitoring agent in 10 minutes.

Windows and UNIX systems

Taking operating system agents as a baseline, then in theory for UNIX and Windows you can install them in 10 minutes. However, when planning, allow time to debug some failures. Ideally, in an 8-hour day you should be able to install 48 error-free agents. If scripts are allowed to run overnight, nearly 150 agents can be installed per day. Even so this is probably not realistic for a couple of reasons.

First, you must spend time debugging problems and due to issues with access to user IDs that have sufficient authority to perform the installation and access to the installation image. In most environments a single person can reasonably aim for 50 agents per day if user access is provided in a timely manner. Even if there are multiple agents to be installed per computer, then this is still a reasonable estimate for installing an operating system agent and a database agent.

Note: If software distribution is being used, then after a package has been built, the number of agents that can be deployed in a given time frame increase.

z/OS systems

If you take CICS, DB2, and z/OS agents as an example, in general allow one day to configure each agent, assuming one CICS region and one DB2 subsystem. For additional CICS regions, this means more CICS table changes. Additional DB2 subsystems are defined in ICAT and you can estimate one hour per subsystem.

Upgrading your environment to allow for additional LPARs is typically quicker because you can use batch mode replication. Allow one day per LPAR.

For more information search for “SMP/E installation on z/OS” or navigation code “1TW10TM3M” at the IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>).

Self-describing monitoring agents

The self-describing agent feature in Tivoli Monitoring makes it possible for new or updated IBM Tivoli Monitoring agents to become operational after installation, without having to perform additional product support installation steps.

Self-describing agents apply version updates to other components automatically without the need to recycle your hub Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, or remote Tivoli Enterprise Monitoring Server. This feature integrates the installation of an agent with the dispersal and installation of associated product support files throughout your IBM Tivoli Monitoring infrastructure.

For more information, see “Self-describing agent installation” on page 394.

Setting up LDAP and single sign-on for your dashboard environment

If you are setting up a dashboard environment with Dashboard Application Services Hub and monitoring dashboard applications such as IBM Infrastructure Management Dashboards for Servers, and you want to control which monitoring resources each of your dashboard users can access, you must set up an LDAP user repository if you do not already have one in your environment. Then configure Dashboard Application Services Hub and your portal servers to use the LDAP repository and enable single sign-on.

If your organization does not have a LDAP server or LDAP skills then allow time to determine which LDAP server to use, such as Tivoli Directory Server or Microsoft Active Directory. Also allocate time for an administrator to become familiar with LDAP concepts and to setup and configure the server and add users and groups to the repository.

Once the LDAP server is setup, allow one day to configure Dashboard Application Services Hub and your portal servers to use the LDAP server and to enable single sign-on. See the Jazz for Service Management Configuration Guide for information on setting up Dashboard Application Services Hub to

use LDAP and single sign-on. See the *IBM Tivoli Monitoring Administrator's Guide* for information on setting up the portal server to use LDAP and single sign-on.

Creating authorization policies

If you plan to use authorization policies to control which managed systems and managed system groups a user or user group can view in monitoring dashboards, allocate a day to become familiar with creating authorization policies and using the `tivcmd` Command Line Interface.

The authorization policies must be created before you can enable authorization policy enforcement in the portal server.

Setting up situation-based monitoring

It is difficult to estimate the time required to set up situation-based monitoring. It is easy to turn on various predefined situations if they are not already running, but not always practical.

Before starting or turning off predefined situations, you need to review them to determine which ones are appropriate for your environment. Moreover, you need to review the default thresholds for the situations that you decide to use and adjust them to suit your site policies and requirements.

Note: Any changes made to predefined situations are overwritten when application support is updated. If you modify predefined situations, save them with a new name, and use them instead of the original predefined situation for your monitoring.

For creating new situations, allow for 20 situations per day, including testing.

You may also want to create managed system groups, which are lists of managed systems to which you want to distribute the same situations.

When creating situations and managed systems lists, it is extremely important to choose a naming convention that is meaningful. A similar naming convention must be used for situations, managed systems lists, workspaces, and queries. "Customizing your environment" on page 99 describes some options when considering the situation and managed system group naming conventions.

Creating workflow policies

You can use workflow policies to perform complex monitoring and automation. Time must be allocated to determine when situations are inadequate to provide complex thresholding and Take Actions.

When you have determined which complex monitoring and automation cannot be accomplished, put proper plans in place to ensure effective workflow policies are created. When planning workflow policies, choose a naming convention that is similar to your convention for naming situations. Plan enough time to adequately develop and test your workflow policies.

Creating workspaces

Creating Tivoli Enterprise Portal workspaces is difficult to estimate. Typically, you do not know exactly what you want to monitor. Be careful not to create something that will be hard to maintain in the future.

Workspaces vary in complexity a great deal, so a graphic view with icons does not take long, but a workspace with links using variables takes quite a while to set up and test. On balance, aim for ten workspaces per day along with associated custom navigator items.

Creating and deploying Agent Builder Agents

Installing Agent Builder agents is almost impossible to estimate because of the varied nature of what the Agent Builder agents can monitor. If you are using user-provided scripts, for instance, then you can

create an application in half an hour. If you have to create your own scripts then that can take days depending on your skill level with scripting and the complexity of the script. Likewise, log file monitoring can be straightforward or complex depending on the format of the log file. Building an SNMP based agent can take a few minutes, building one that uses the Java Custom API to interact with custom applications can take days or even longer. A single agent can include any combination of providers and between one and more than one hundred attribute groups.

Adjust the time estimates on a case-by-case basis. Be very careful to commit the agents before installing them into production environments. This tells the Agent Builder to enforce the rules that allow a clean upgrade path between agent versions.

Transferring skills

Much of the skills transfer can occur during the installation and configuration process if you are working closely with key staff. You can still factor in an additional three days to cover the same subjects as in the IBM Tivoli Monitoring Administration course.

Other skills transfer can be estimated at one day per agent type. For some of the z/OS-based agents this can be a bit longer because of the two different 3270 mainframe and the portal client interfaces, so allow two days for CICS and DB2 agents.

These time estimates correspond roughly to the number of days it takes to deliver formal training.

Scheduling the initial deployment

The initial deployment of the Tivoli Monitoring environment is the most time-consuming. The majority of the customization is performed during or immediately after the initial deployment.

Customization includes modifying workspaces, modifying situation thresholds, building managed systems lists, and defining Take Action commands.

Scheduling for fix packs

In a typical user environment, allocate one day to upgrade the Tivoli Monitoring infrastructure components (hub monitoring server, portal server, Warehouse Proxy Agent, Summarization and Pruning Agent). In a multi-hub environment, allow one day for each hub. Then, time must be allocated to upgrade the agents.

Typically, the schedule is driven largely by Change Control windows and not by the time it takes to deploy the fix pack to the agents. However, for planning purposes use the following time estimates.

These times assume that there is adequate network bandwidth between the remote monitoring server deployment depot and the agents. In environments with slow network links between the remote monitoring server and the agents, evaluate the size of the fix pack files to calculate the rate of transfer. Because each fix pack is a different size and each network has unique characteristics, calculate this for your environment before starting the upgrade.

Performing agent upgrades in parallel can be done with the `itmpatchagents` tool. For environments with adequate network bandwidth and two people performing upgrades in parallel, plan to upgrade approximately 500 agents per day.

Finally, following the installation of a fix pack, plan time to test your environment. For Tivoli Monitoring infrastructure fix packs, spend two person days thoroughly testing your environment. For application agents, there is little risk that the Tivoli Monitoring infrastructure components were affected by the fix pack. Therefore, no more than one person day need be allocated to test the environment following the installation of an application fix pack.

Staffing

The table in this section lists the staffing and time estimates for various Tivoli Monitoring tasks.

Table 9. Staffing estimates

Tivoli Monitoring tasks	Hours required	Number of people required	Skills required
Discussing the project plan including implementation and support in place by IBM	16	2	Medium
Downloading the correct media for installation	8 hours	1	Medium
Completing the hardware check	4 hours	1	Medium
Creating the warehouse database with sufficient space, user IDs and tuning parameters	4	1	High (DBA)
Configuring firewall for ports	8	1	High (Network Team)
Installing core components (portal server, monitoring server, Warehouse Proxy Agent, and Summarization and Pruning Agent) ¹	20	2	High
Disabling the default situations	2	1	Medium
Creating the managed systems list	8	1	Medium
Creating the custom situations	32	2	Medium
Deploying first 50 mix of Windows OS, UNIX and Linux OS and different application agents	40	1	High
Deploying situations to appropriate managed systems lists	8	1	Medium
Verifying all agents and core components are active and working correctly	40	2	High
Verifying events flowing to Tivoli Netcool/OMNIBus ObjectServer or Tivoli Enterprise Console server	4	1	Medium
Backing up the environment (core components)	4	1	High
Deploying agents in chunks of 100 agents, depending upon the requirement. Sizing includes adding computers to the managed systems lists.	24	2	Medium
Enabling only required attribute groups for warehousing as mentioned in the project plan	8	1	Medium
Verifying warehouse data is flowing from agents to warehouse database	8	1	High
Performing frequent health checks of the environment, more frequent in first few months and then later less frequent checks	4	1	High
Configuring Performance Analyzer Agent	.5	1	Medium
Adding application support to Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server and Tivoli Enterprise Portal Client	.5	1	Medium
Reconfiguring Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, and Tivoli Enterprise Portal Client	.5	1	Medium
Configuring historical data collection	2	1	Medium
Moving the Tivoli Performance Analyzer configuration database ²	4	1	Medium
Installing Tivoli Performance Analyzer domain support in console mode or in GUI mode	2	1	Medium
Installing reports	1	1	Medium

Table 9. Staffing estimates (continued)

Tivoli Monitoring tasks	Hours required	Number of people required	Skills required
Inspection, Validation, and Enabling or Disabling of the default tasks to suit the environments and requirements	2	1	Medium
Customizing the default tasks to suit the environment and requirements	2	1	Medium
Creating custom tasks	2	1	Medium
Note: 1. Installing core components depends on the size of the environment and the complexity regarding multiple Warehouse Proxy Agents. 2. In some situations you might want to move the Tivoli Data Warehouse database to a different computer or Relational Data Base Management System. The migration procedure described in Tivoli Data Warehouse manuals does not cover Performance Analyzer configuration data. For more information, see <i>Moving the Tivoli Performance Analyzer configuration database</i> in the IBM Tivoli Monitoring information center.			

The steps outlined in Chapter 4, “Deployment phase,” on page 97 will assist you with your project planning and ensure that the deployment phase goes smoothly.

Chapter 3. Hardware and software requirements

The following sections provide specific information about the memory, software, and hardware requirements for installing IBM Tivoli Monitoring.

Note: This section does not show agent-specific requirements (such as supported application levels or any hardware requirements unique to a certain agent). For this information, see the user's guide (for agents on distributed systems) or configuration guide (for agents on mainframe systems) for the agent that you are installing.

You can perform prerequisite scanning for agents before carrying out an installation. The two mechanisms available are a manually executed stand-alone prerequisite scanner and a remote prerequisite scanner facility that extends the capabilities of IBM Tivoli Monitoring's remote deployment component. For more information, see "Prerequisite Scanner for IBM Tivoli Monitoring agents" on page 48.

This chapter contains the latest Hardware and Software requirements information available at the time of publication. IBM Tivoli Monitoring Hardware and Software requirements information is also available in the Tivoli Monitoring Wiki. You can download the PDF from the Application Performance Management community **Files** section.

Supported operating systems

The following tables show which operating systems are supported for the different IBM Tivoli Monitoring components: monitoring server, automation server, portal server, portal client, monitoring agent, Warehouse Proxy, Warehouse Proxy Summarization and Pruning Agent, and Tivoli Performance Analyzer.

Note:

1. In the following tables, an X indicates a supported component on the platform. If the platform being discussed supports 64 bit applications but only a 32 bit version of the Tivoli Monitoring component is supported on that platform, the X is replaced by "32 bit."
2. The tables in this section show support information for monitoring components. For prerequisite information about applications that are installed into the Dashboard Application Services Hub, such as the Tivoli Authorization Policy Server, and the IBM Infrastructure Management Dashboards for Servers, see the Hardware and software requirements in the Jazz for Service Management information center at http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install_prereqs.html. The tivcmd Command Line Interface for Authorization Policy is also supported on the same operating systems as IBM Dashboard Application Services Hub. The tivcmd Command Line Interface for Authorization Policy is also supported on the same 32 bit and 64 bit Windows operating systems, and the 32 bit and 64 bit Linux Intel based operating systems, that are supported by the IBM Tivoli Monitoring portal client as shown in the tables below. However, on a 64 bit zSeries Linux system, you must install the 32 bit version of the tivcmd CLI.

Table 10 shows the support for monitoring components on Windows computers.

Table 10. Supported Windows operating systems

Operating system	Monitoring server and automation server	Portal server ^{1 4}	Portal client ²	OS monitoring agent ³	Warehouse Proxy Agent	Summarization and Pruning Agent	Tivoli Performance Analyzer
Windows Server 2003 Standard Edition R2 on Intel x86-32 (32 bit)			X	X			

Table 10. Supported Windows operating systems (continued)

Operating system	Monitoring server and automation server	Portal server ^{1 4}	Portal client ²	OS monitoring agent ³	Warehouse Proxy Agent	Summarization and Pruning Agent	Tivoli Performance Analyzer
Windows Server 2003 Standard Edition R2 on x86-64 (64 bit)			X	X			
Windows Server 2003 Enterprise Edition R2 on Intel x86-32 (32 bit)			X	X			
Windows Server 2003 Enterprise Edition R2 on x86-64 (64 bit)			X	X			
Windows Server 2003 Datacenter Edition R2 on Intel x86-32 (32 bit)			X	X			
Windows Server 2003 Datacenter Edition R2 on Intel x86-64 (64 bit)			X	X			
Windows Server 2008 Standard Edition on Intel x86-32 (32 bit) ²	X	X	X	X	X	X	X
Windows Server 2008 Standard Edition on Intel x86-64 (64 bit) ²	X	X	X	X	X	32 bit	32 bit
Windows Server 2008 Enterprise Edition on Intel x86-32 (32 bit) ²	X	X	X	X	X	X	X
Windows Server 2008 Enterprise Edition on Intel x86-64 (64 bit) ²	X	X	X	X	X	32 bit	32 bit
Windows Server 2008 Datacenter Edition on Intel x86-32 (32 bit) ²	X	X	X	X	X	X	X
Windows Server 2008 Datacenter Edition on Intel x86-64 (64 bit) ²	X	X	X	X	X	32 bit	32 bit
Windows 7 Enterprise Edition on Intel x86-32 (32 bit)			X	X			
Windows 7 Enterprise Edition on Intel x86-64 (64 bit)			X	X			
Windows 7 Professional Edition on Intel x86-32 (32 bit)			X	X			
Windows 7 Professional Edition on Intel x86-64 (64 bit)			X	X			
Windows 7 Ultimate Edition on Intel x86-32 (32 bit)			X	X			
Windows 7 Ultimate Edition on Intel x86-64 (64 bit)			X	X			
Windows 8 Pro on Intel x86-32 (32 bit)			X	X			

Table 10. Supported Windows operating systems (continued)

Operating system	Monitoring server and automation server	Portal server ^{1 4}	Portal client ²	OS monitoring agent ³	Warehouse Proxy Agent	Summarization and Pruning Agent	Tivoli Performance Analyzer
Windows 8 Pro on Intel x86-64 (64 bit)			X	X			
Windows 8 Enterprise on Intel x86-32 (32 bit)			X	X			
Windows 8 Enterprise on Intel x86-64 (64 bit)			X	X			
Windows Server 2008 R2 Standard Edition on Intel x86-64 (64 bit) ²	X	X	X	X	X	32 bit	32 bit
Windows Server 2008 R2 Enterprise Edition on Intel x86-64 (64 bit) ²	X	X	X	X	X	32 bit	32 bit
Windows Server 2008 R2 Datacenter Edition on Intel x86-64 (64 bit) ²	X	X	X	X	X	32 bit	32 bit
Windows Server 2008 R2 Server Core on Intel x86-64 (64 bit)				X			
Windows Server 2012 Essentials Edition on Intel x86-64 (64 bit)	X	X	X	X	X	32 bit	32 bit
Windows Server 2012 Standard edition on Intel x86-64 (64 bit)	X	X	X	X	X	32 bit	32 bit
Windows Server 2012 Datacenter edition on Intel x86-64 (64 bit)	X	X	X	X	X	32 bit	32 bit
Note: <ol style="list-style-type: none"> 1. The Tivoli Enterprise Portal Server runs an instance of embedded WebSphere Application Server V8, so the system must also meet the minimum requirements for WebSphere. For information on the minimum service pack level required to install and run the embedded WAS in Tivoli Enterprise Portal Server, see http://www-01.ibm.com/support/docview.wss?uid=swg27021246. 2. The Tivoli Enterprise Portal desktop client is supported on marked platforms. The Tivoli Enterprise Portal browser client can run from the marked platforms in the browser versions documented in "Supported browser versions" on page 90. The Tivoli Enterprise Portal Java webstart client can run from the marked platforms. 3. The OS monitoring agent column indicates the platforms on which an operating system monitoring agent is supported. This column does not indicate that any agent runs on any operating system. For example, to monitor a Linux computer, you must use a Linux monitoring agent, not a Windows monitoring agent. For information about the operating systems supported for non-OS agents, see the documentation for the specific agents you are using in your environment. 4. Beginning with V6.3, Microsoft Visual C++ 2008 Redistributable Package (32 bit and 64 bit) is a prerequisite for the Tivoli Enterprise Portal Server on Windows systems (and the System Monitor Agent on Windows). For more information go to the Microsoft Download Center at http://www.microsoft.com/en-us/download/default.aspx. 							

Table 11 shows the support for monitoring components on UNIX (non-Linux), IBM i, and z/OS computers.

Table 11. Supported UNIX, IBM i, and z/OS operating systems

Operating system	Monitoring server and automation server	Portal server	Portal client	OS monitoring agent ^{1, 2}	Warehouse Proxy Agent ³	Summarization and Pruning Agent	Tivoli Performance Analyzer
AIX V6.1 (64 bit) ⁴	X	X ¹²		X	X	X	32 bit

Table 11. Supported UNIX, IBM i, and z/OS operating systems (continued)

Operating system	Monitoring server and automation server	Portal server	Portal client	OS monitoring agent ^{1, 2}	Warehouse Proxy Agent ³	Summarization and Pruning Agent	Tivoli Performance Analyzer
AIX V7.1 (64 bit) ^{4, 13}	X	X		X	X	X	32 bit
Solaris V10 (SPARC) (32/64 bit) ⁵	32 bit			X	32 bit	32 bit	32 bit
Solaris V10 (Intel x86-64) (64 bit) ⁵	32 bit			X			
Solaris V11 (SPARC) (32/64 bit)	32 bit			X	32 bit	32 bit	32 bit
Solaris V11 (Intel x86-64) (64 bit)	32 bit			X			
Solaris Zones (SPARC) (32/64 bit) ⁸	32 bit			X	32 bit	32 bit	32 bit
Solaris Zones (Intel x86-64) (64 bit) ⁸	32 bit			X			
HP-UX 11i v2 (64 bit) on PA-RISC ^{10, 11}				X			
HP-UX 11i v3 (64 bit) on PA-RISC ^{10, 11}				X			
HP-UX 11i v2 on Integrity (IA64) ^{9, 10, 11}				X			
HP-UX 11i v3 on Integrity (IA64) ^{9, 10, 11}				X			
i5/OS 5.4 (64 bit) ⁶				X			
IBM i 6.1 (64 bit) ⁶				X			
IBM i 7.1 (64 bit) ⁶				X			
z/OS 1.12 (31/64 bit) ^{7, 8}	31 bit ¹⁴			X			
z/OS 1.13 (31/64 bit) ^{7, 8}	31 bit ¹⁴			X			

Note:

1. The **OS monitoring agent** column indicates the platforms on which an operating system monitoring agent is supported. This column does not indicate that any agent runs on any operating system. For example, to monitor a Linux computer, you must use a Linux monitoring agent, not a Windows monitoring agent.
For information about the operating systems supported for non-OS agents, see the documentation for the specific agents you are using in your environment.
2. If you are installing the OMEGAMON XE for Messaging agent on a 64-bit operating system, you must install the 32-bit version of the agent framework.
3. Configuration of the Warehouse Proxy Agent requires an X Window System (also known as the X11 GUI) on the computer where you are configuring it. Alternatively, you can run the following command to use an X terminal emulation program (such as Cygwin) that is running on another computer:

```
export DISPLAY=my_windows_pc_IP_addr:0.0
```


where *my_windows_pc_IP_addr* is the IP address of a computer that is running an X terminal emulation program.
4. Version 11.1 of the AIX XL C/C++ runtime must be installed. See the following Web site for installation instructions:
<http://www-01.ibm.com/support/docview.wss?uid=swg24031426>.
5. There are some limitations on installing into Solaris 10 when zones are configured. See “Installing into Solaris zones” on page 135.
6. SNMP version 3 is not supported on IBM i.
7. For information about installing the monitoring server on z/OS, see the program directory that comes with that product.
8. The OS monitoring agent for z/OS computers is part of the IBM Tivoli OMEGAMON for z/OS product.

Table 11. Supported UNIX, IBM i, and z/OS operating systems (continued)

Operating system	Monitoring server and automation server	Portal server	Portal client	OS monitoring agent ^{1, 2}	Warehouse Proxy Agent ³	Summarization and Pruning Agent	Tivoli Performance Analyzer
<p>9. You cannot upgrade either the OS or Log Alert agents that you currently have running on an HP-UX 11i v2 (B.11.23) on an Integrity (IA64) computer in PA-RISC mode prior to v6.1 fix pack 4. Fix packs prior to fix pack 4 did not run in native 64-bit mode by default. You must first uninstall the agent if the version is previous to the V6.1 fix pack 4 version.</p> <p>10. The 32-bit kernel still requires a 64-bit processor. Ensure that any HP-UX managed system is based on PA-RISC2 architecture. From the native kernel mode (for example, 64 bit if the system is 64-bit based), run the following command: file /stand/vmunix. This returns the native architecture type. For example: /stand/vmunix: PA-RISC1.1 executable -not stripped. Verify the architecture is at least PA-RISC2.</p> <p>11. Prerequisite patches for HP-UX R11: PHSS_26946, PHSS_33033, and PHCO_34275 or later.</p> <p>12. The Tivoli Enterprise Portal server runs an instance of embedded WebSphere Application Server V8, so users of AIX must also meet the minimum requirements for WebSphere. For information on the minimum service pack level required to install and run the embedded WAS in Tivoli Enterprise Portal Server, see http://www-01.ibm.com/support/docview.wss?uid=swg27021246.</p> <p>13. On AIX 7.1 TL1, Service Pack (SP) 2 or later is required for all components. If you cannot upgrade to SP 2, you can download the fix for APAR IV09585. For more information, see the <i>IBM Tivoli Monitoring Troubleshooting Guide</i>. You must restart your system after you have upgraded to SP 2 or downloaded the fix for APAR IV09585.</p> <p>14. The Tivoli Enterprise Automation Server is not supported on z/OS.</p>							

Table 12 shows the monitoring components supported on Linux operating systems.

Table 12. Supported Linux operating systems

Operating system	Monitoring server and automation server ⁸	Portal server ^{7 8}	Portal client ^{1, 8}	OS monitoring agent ^{2, 3, 8}	Warehouse Proxy Agent ⁸	Summarization and Pruning Agent ⁸	Tivoli Performance Analyzer ⁸
Asianux 3.0 for Intel x86-32 (32 bit) ⁶				X			
Asianux 3.0 on Intel x86-64 (64 bit) ⁶				X			
Red Flag 5.0 for Intel x86-32 (32 bit) ⁶				X			
RedHat Enterprise Linux 5 Intel x86-32 (32 bit) ^{4 6}	X	X ⁹	X	X	X	X	X
RedHat Enterprise Linux 5 on Intel x86-64 (64 bit) ^{4 6}	X	X ⁹	X	X	X	X	32 bit
RedHat Enterprise Linux 5 on iSeries and pSeries (64 bit) ⁴				X			
RedHat Enterprise Linux 5 on zSeries (64 bit) ^{4 6}	X	X ^{6 9}		X	X	X	
RedHat Enterprise Linux 6 Intel x86-32 (32 bit) ^{4 6}	X	X	X	X	X	X	X
RedHat Enterprise Linux 6 on Intel x86-64 (64 bit) ^{4 6}	X	X	X	X	X	X	32 bit
RedHat Enterprise Linux 6 on iSeries and pSeries (64 bit) ⁵				X			

Table 12. Supported Linux operating systems (continued)

Operating system	Monitoring server and automation server ⁸	Portal server ^{7 8}	Portal client ^{1, 8}	OS monitoring agent ^{2, 3, 8}	Warehouse Proxy Agent ⁸	Summarization and Pruning Agent ⁸	Tivoli Performance Analyzer ⁸
RedHat Enterprise Linux 6 on zSeries (64 bit) ^{4 5}	X	X ⁶		X	X	X	
SuSE Linux Enterprise Server 10 Intel x86-32 (32 bit) ⁶	X	X ¹⁰	X	X	X	X	X
SuSE Linux Enterprise Server 10 on Intel x86-64 (64 bit) ⁶	X	X ¹⁰	X	X	X	X	X
SuSE Linux Enterprise Server 10 for iSeries and pSeries (64 bit)				X			
SuSE Linux Enterprise Server 10 for zSeries (64 bit) ⁶	X	X ^{5 10}		X	X	X	
SuSE Linux Enterprise Server 11 on Intel x86-32 (32 bit) ⁶	X	X	X	X	X	X	X
SuSE Linux Enterprise Server 11 on Intel x86-64 (64 bit) ^{4 6}	X	X	X	X	X	X	32 bit
SuSE Linux Enterprise Server 11 for iSeries and pSeries (64 bit)				X			
SuSE Linux Enterprise Server 11 for zSeries (64 bit) ⁶	X	X ⁵		X	X	X	
VMWare ESX Server 3.0.1 on Intel x86-32 (32 bit)				native Linux OS agent			
VMWare ESX Server 3.0.1 on Intel x86-64 (64 bit)				native Linux OS agent			
VMWare ESX Server 3.5 on Intel x86-32 (32 bit)				native Linux OS agent			
VMWare ESX Server 3.5 on Intel x86-64 (64 bit)				native Linux OS agent			
VMWare ESX Server 4.0 on Intel x86-32 (32 bit)				native Linux OS agent			
VMWare ESX Server 4.0 on Intel x86-64 (64 bit)				native Linux OS agent			

Table 12. Supported Linux operating systems (continued)

Operating system	Monitoring server and automation server ⁸	Portal server ^{7 8}	Portal client ^{1, 8}	OS monitoring agent ^{2, 3, 8}	Warehouse Proxy Agent ⁸	Summarization and Pruning Agent ⁸	Tivoli Performance Analyzer ⁸
<p>Note:</p> <ol style="list-style-type: none"> The Tivoli Enterprise Portal desktop client is supported on marked platforms. The Tivoli Enterprise Portal browser client can run from the marked platforms in the browser versions documented in "Supported browser versions" on page 90. The Tivoli Enterprise Portal Java webstart client can run from the marked platforms. The OS monitoring agent column indicates the platforms on which an agent is supported. This column does not indicate that any agent runs on any operating system. For example, to monitor a Linux computer, you must use a Linux monitoring agent, not a Windows monitoring agent. An "X" symbol in the column indicates that an operating system agent is available for the specific operating system that is named in the row where the "X" is located. The Linux OS Monitoring Agent requires the installation of the latest versions of the following libraries: libstdc++ libgcc compat-libstdc++ These libraries are available on the Linux operating system installation media and Service Packs. Each library can have multiple packages, and each should be installed. If you are on a 64 bit system you may require both the 32 bit and 64 bit versions of these libraries, depending on which components you install. Several Linux operating systems now enable SELinux (security-enhanced Linux) by default, which interferes with the installation, configuration, and operation of IBM Tivoli Monitoring. To ensure the proper operation of the V6.1.x and V6.2.x, and V6.3 releases, the SELinux setting must be either disabled, or changed from enforcing mode to permissive mode when enabled. When the permissive mode is chosen, the system log will contain entries regarding which Tivoli Monitoring binaries have triggered the SELinux security condition. However, under the permissive mode, these entries are for auditing purpose only, and will operate normally. For Linux on System z[®], a 64 bit Tivoli Enterprise Portal Server requires running a 64 bit DB2 instance under Linux. The Tivoli Enterprise Portal server runs an instance of embedded WebSphere Application Server V8, so the system must also meet the minimum requirements for WebSphere. For information on the minimum service pack level required to install and run the embedded WAS in Tivoli Enterprise Portal Server, see http://www-01.ibm.com/support/docview.wss?uid=swg27021246. When installing as a nonroot user, 32 bit libpam is required for the Tivoli Enterprise Portal Server on a 64 bit Intel Linux installation. The following libraries are required to run Manage Tivoli Enterprise Monitoring Services (CandleManage) on Linux systems: <ul style="list-style-type: none"> libXcursor libXfixes libXi RedHat Enterprise Linux V5.6 or later is required for the Tivoli Enterprise Portal Server. SuSE Linux Enterprise Server 10.3 or later is required for the Tivoli Enterprise Portal Server. 							

The following table lists the operating system patches required for the IBM Global Security Toolkit (GSKit), which is used to provide security between monitoring components. GSKit is installed automatically when you install Tivoli Management Services components and the tivcmd Command Line Interface.

Note: In 64-bit environments, only the 64-bit GSKit is installed by default. The 32-bit GSKit is no longer installed by default. The 32-bit GSKit is only installed if a 32-bit component that requires 32-bit GSKit is installed into the 64-bit environment.

Table 13. Operating system requirements for IBM GSKit

Operating system	Patch required
Solaris V9	111711-08
Solaris V10	none
Solaris V11	none
HP-UX V11i	PHSS_26946, PHSS_33033, PHCO_34275

Table 13. Operating system requirements for IBM GSKit (continued)

Operating system	Patch required
AIX V6.x	xlC.aix61.rte 9.0
Windows Server 2003	none
Red Hat Enterprise Linux 6 x86 (32 bit)	ksh-20100621-2.el6.i686.rpm
Red Hat Enterprise Linux 6 x86-64 (64 bit)	ksh-20091224-1.el6.x86_64.rpm glibc-2.12-1.7.el6.i686.rpm libgcc-4.4.4-13.el6.i686.rpm nss-softokn-freebl-3.12.7-1.1.el6.i686.rpm
SuSE Linux Enterprise Server 10 Intel	none
SuSE Linux Enterprise Server 11 Intel	none
Linux on System z	31 bit (s390) glibc and libgcc must be installed on a 64 bit system if you are installing 31 bit components.

Supported databases for Tivoli Enterprise Portal Server

IBM Tivoli Monitoring includes an embedded version of the Apache Derby database server for default use as the Tivoli Enterprise Portal Server database. The database is installed when installing the portal server, and it runs within the portal server's Java virtual machine.

The embedded Derby database is intended for use only by small to medium-size businesses. It is not intended for large IBM Tivoli Monitoring environments or for sites running a large number of Tivoli Enterprise Portal clients. Derby has been tested with up to 20 portal clients, and it does require higher CPU and memory usage than DB2 for Linux, UNIX, and Windows.

If the embedded Derby database is not suitable for your environment, you can use one of the supported databases in Table 14. Note that the database and the portal server must be installed on the same computer.

Note: IBM Tivoli Monitoring V6.3 includes DB2 Enterprise Server Edition 10.1 for use with the portal server and the Tivoli Data Warehouse.

Table 14. Supported databases for the portal server

Portal server operating system	Portal server database ("TEPS") ^{1, 2}	
	IBM DB2 for Linux, UNIX, and Windows	MS SQL Server
AIX	<ul style="list-style-type: none"> • V9.7 and fix packs • V9.8 and fix packs • V10.1 and fix packs 	
Linux	<ul style="list-style-type: none"> • V9.7 and fix packs • V9.8 and fix packs • V10.1 and fix packs 	
Windows	<ul style="list-style-type: none"> • V9.7 and fix packs • V9.8 and fix packs • V10.1 and fix packs 	<ul style="list-style-type: none"> • MS SQL Server 2008 • MS SQL Server 2008 R2 • MS SQL Server 2012

Table 14. Supported databases for the portal server (continued)

Portal server operating system	Portal server database ("TEPS") ^{1, 2}	
	IBM DB2 for Linux, UNIX, and Windows	MS SQL Server
Note: <ol style="list-style-type: none"> "TEPS" is the default database name for the database used by the portal server. Your portal server database must be located on the computer where the portal server is installed. On Linux, the portal server database must be installed with the operating system language set to UTF-8. If you transition from one supported database system to another for the Tivoli Enterprise Portal Server, your existing portal server data is not copied from the first system to the new one. The Tivoli Enterprise Portal Server automatically prunes closed events out of the database after 24 hours, to prevent long-term database growth. 		

Supported databases for Tivoli Data Warehouse

The following tables show the supported databases for the Tivoli Data Warehouse.

Table 15 shows the supported databases for the Tivoli Data Warehouse. If you run the database for the Tivoli Enterprise Portal Server and the database for the warehouse in the same instance of IBM DB2 for Linux, UNIX, and Windows, you must follow the support requirements in Table 14 on page 82.

Note: IBM Tivoli Monitoring V6.3 includes DB2 Enterprise Server Edition 10.1 for use with the portal server and the Tivoli Data Warehouse.

Table 15. Supported databases for the Tivoli Data Warehouse

Tivoli Data Warehouse database ("WAREHOUS") ¹			
IBM DB2 for Linux, UNIX, and Windows	IBM DB2 on z/OS ⁵	MS SQL Server	Oracle
Supported versions: <ul style="list-style-type: none"> V9.7 and fix packs V10.1 and fix packs 	Supported versions: version 9.1 or higher. Support applies to any Windows, Linux, or UNIX platform that can run the Warehouse Proxy Agent. DB2 Connect [™] Server Edition is also required on the workstation.	Supported versions: <ul style="list-style-type: none"> MS SQL Server 2008 Enterprise Edition MS SQL Server 2008 R2 Enterprise Edition MS SQL Server 2012 Enterprise Edition 	Supported versions: <ul style="list-style-type: none"> 11g Release 1 11g Release 2
Note: <ol style="list-style-type: none"> "WAREHOUS" is the default database name for the database used by Tivoli Data Warehouse. Support is for 32-bit or 64-bit databases. Your Tivoli Data Warehouse database can be located on the same computer as your monitoring server or on a remote computer. See the Oracle company support Web site (http://www.oracle.com/index.html) for information about installing and configuring Oracle on Solaris V10. DB2 on z/OS is not supported for the Tivoli Enterprise Portal Server. 			

Table 16 on page 84 shows the supported databases for the Tivoli Data Warehouse database compression feature.

Table 16. Supported databases for Tivoli Data Warehouse database compression

Tivoli Data Warehouse database compression		
IBM DB2	MS SQL Server	Oracle
V9.7 or higher, with the IBM DB2 Storage Optimization feature.	MS SQL Server 2008 Enterprise Edition or higher.	11g Release 1 or higher, with the Oracle Advanced Compression feature.

Note: DB2 on z/OS is not supported for this feature.

Table 17 shows the supported databases for the Tivoli Data Warehouse database range partitioning feature.

Table 17. Supported databases for Tivoli Data Warehouse database range partitioning

Tivoli Data Warehouse range partitioning		
IBM DB2	IBM DB2 on z/OS	Oracle
Enterprise Edition 9.7 or higher.	DB2 for z/OS 9.1 or higher.	11g Release 1 or higher, with the partitioning option.

Required hardware for distributed systems

The following sections describe the processor, disk, memory, and other hardware requirements for the IBM Tivoli Monitoring infrastructure components on distributed systems.

A *distributed system* is defined here as any hardware that is not zSeries.

Following is a list of the IBM Tivoli Monitoring components covered in this section:

- Hub monitoring server
- Remote monitoring server
- Tivoli Enterprise Monitoring Automation Server
- Portal server
- Portal client
- IBM Dashboard Application Services Hub with IBM Infrastructure Management Dashboards for Servers
- Tivoli Data Warehouse
- Warehouse Proxy Agent
- Summarization and Pruning Agent

Processor requirements

For best performance, processor speeds should be as fast as possible. Choosing faster processors should result in improved response time, greater throughput, and lower CPU utilization. For Windows or Linux on AMD and Intel systems, an x64 processor is recommended.

Multiprocessor systems are generally recommended for the IBM Tivoli Infrastructure components. The infrastructure components (monitoring server, portal server, portal client) run as multithreaded processes and are able to run threads concurrently across multiple processors if they are available. CPU utilization for most components is bursty, and steady-state CPU utilization is expected to be low in most cases. For components supporting large environments, using multiprocessor systems can improve throughput.

When the infrastructure components are installed in a virtualized environment (PowerVM, VMWare, and so on), virtual machines with multiple logical processors are recommended, but dedicated processors are

not usually required. Capping the CPU usage of these virtual machines is not generally recommended, so that processing for bursts of activity can utilize spare capacity on the system.

When multiple infrastructure components are installed in a single system or virtual machine, additional processors might be needed for best performance.

For large environments, consider using four or more processors for the Tivoli Data Warehouse database server and the Summarization and Pruning Agent. Although it is not required, the Summarization and Pruning Agent can run on a machine separate from the Tivoli Data Warehouse database server. In this case, ensure that a high-speed network connection exists (100 Mbps or faster) between the Summarization and Pruning Agent and the database server.

Memory and disk requirements

The following table shows estimated memory and disk storage for IBM Tivoli Monitoring components on distributed systems. The estimated memory usage for each component does not include the operating system and any concurrently running applications. To estimate the memory requirements for a machine, add the memory requirements for each component that will run on the machine, and add at least 1 GB for the operating system and to allow memory for filesystem caching.

Add the sizings for individual components to calculate a total for more than one component installed on the same computer. For example, if the hub monitoring server and portal server are installed on the same machine, the base component memory requirement is 1600 MB with an additional 220 MB per 1000 managed systems. For an environment with 2000 managed systems, the combined component memory requirement is $1600 + 2 * 220 = 2040$ MB. The machine should be configured with at least 4 GB RAM (2040 MB + 1 GB for the OS and filesystem caching + extra memory for dynamic requests and some situation evaluation on the hub monitoring server).

Table 18. Estimated memory and disk storage for IBM Tivoli Monitoring components on distributed systems

Component	Memory requirements ¹		Disk storage requirements ²
	Base memory requirement	Additional memory requirements	
Tivoli Enterprise Monitoring Server	100 MB	120 MB per 1000 managed systems in environment. Extra memory for situations evaluated at monitoring server. ⁷ Extra memory for large dynamic data requests. ⁸	1.5 GB ^{3 5}
Remote Tivoli Enterprise Monitoring Server	100 MB	120 MB per 100 managed systems connecting through the server. Extra memory for situations evaluated at monitoring server.	1.0 GB xLinux ^{3 5} 1.5 GB All other systems ^{3 5}
Tivoli Enterprise Monitoring Automation Server	100 MB	70 MB per 1000 managed systems in environment.	100 MB this does not include all of the additional Jazz for Service Management components ⁹

Table 18. Estimated memory and disk storage for IBM Tivoli Monitoring components on distributed systems (continued)

Component	Memory requirements ¹		Disk storage requirements ²
	Base memory requirement	Additional memory requirements	
Tivoli Enterprise Portal Server	1500 MB ⁶	100 MB per 1000 managed systems in environment (for AIX systems, use 150 MB per 1000 managed systems). ⁴ Extra 800 MB if the dashboard data provider is enabled. Extra memory for large dynamic data requests. ⁸	4.0 GB plus an additional 1.2 GB in your computer's temporary directory to install the eWAS server and the Eclipse Help Server ^{5 6}
Portal client (browser or desktop)	400 MB	Extra 200 MB for large environments.	150 MB
Tivoli Data Warehouse	4 GB depending on database configuration parameters	8 - 16 GB depending on database configuration parameters	See "Estimating the required size of your database" on page 510.
Warehouse Proxy Agent	200 MB		150 MB
Summarization and Pruning Agent	500 MB		150 MB
Tivoli Performance Analyzer	200 MB	Add 1.5 GB for large environments.	Depends on Tivoli Data Warehouse configuration for Tivoli Performance Analyzer attribute groups
Dashboard Application Services Hub server (for use with IBM Infrastructure Management Dashboards for Servers)	1000 MB		3.0 GB ⁹

Table 18. Estimated memory and disk storage for IBM Tivoli Monitoring components on distributed systems (continued)

Component	Memory requirements ¹		Disk storage requirements ²
	Base memory requirement	Additional memory requirements	

Note:

- The memory and disk sizings shown in this table are the amounts required for the individual component beyond the needs of the operating system and any concurrently running applications. For the total system memory required by small, medium-size, and large environments, see “Sizing your Tivoli Monitoring hardware” on page 34.
- The disk storage estimates apply to any size monitoring environment and are considered high estimates. The size of log files affect the amount of storage required.
- The disk storage requirements for the hub and remote monitoring servers do not include storage for the agent depot, which can require an additional 1 GB or more.
- The memory requirement for the portal server does not include database processes for the portal server database. DB2 memory usage varies from machine to machine, and can be controlled through various tuning parameters, including bufferpool size. DB2 systems frequently have self-tuning enabled to control the bufferpool size. On systems with a large amount of memory this can result in bufferpools that are much larger than necessary for the TEPS database. A good starting bufferpool size for the TEPS database is 10,000 4K pages (about 40MB). Manual control of the TEPS bufferpool size is recommended to reduce memory usage. You can use the DB2 ALTER command to change the buffer pool size. For example:

CONNECT TO TEPS;
ALTER BUFFERPOOL IBMDEFAULTBP IMMEDIATE SIZE 10000;
CONNECT RESET;
- Tivoli Monitoring components require additional disk storage when the self-describing agent feature is enabled. The self-describing agent disk storage usage sizes should be multiplied by the number of products in your environment that are expected to perform self-describing agent installations. This estimate is only applicable when the self-describing agent feature is enabled. The disk storage is used by the directory indicated in the Tivoli Enterprise Monitoring Server variable TEMS_MANIFEST_PATH, or the Tivoli Enterprise Portal Server variable TEPS_MANIFEST_PATH.
 - The hub monitoring server requires an estimated 3 to 5 MB of additional disk storage for each product that is performing self-describing agent installations. WebSphere and OMEGAMON agents require more storage, so use the 5 MB per product estimate if you are installing those agents. Otherwise, you can use the 3 MB average size per product.
 - The remote monitoring server requires an estimated 2 MB of additional disk storage for each product that is performing self-describing agent installations.
 - The portal server requires an estimated 3 to 6 MB of additional disk storage for each product that is performing self-describing agent installations since it maintains the new and previous version of an agent product's application support. WebSphere and OMEGAMON agents require more storage, so use the 6 MB per product estimate if you are installing those agents. Otherwise, you can use the 3 MB average size per product.
- On AIX only: IBM Tivoli Monitoring V6.3 will install IBM Installation Manager on the system to install the IBM HTTP Server required by the Tivoli Enterprise Portal Server. If IBM Installation Manager is not already installed on this system, the installation requires 300 MB in the /usr directory.
- Use the following formula to estimate the additional memory required for a situation evaluated at the monitoring server: 2 * number of agents running the situation * number of row instances evaluated per situation interval * row size for attribute group. The row size for an attribute group is available in the agent user's guide. As an example, assume an environment has 3000 Linux agents, each running a situation that uses the COUNT function to calculate the number of disks reported by the Linux_Disk attribute group. Assume also that the agents report on average of 15 rows for Linux_Disk, and the row size for Linux_Disk is 250 bytes. The additional memory required at the monitoring server for this situation would be: 2 * 3000 * 15 * 250 = 22,500,000 bytes.
- Requests to the monitoring server for a large amount of data can increase the dynamic memory usage significantly. For example, if a user request results in 10,000 rows of Linux_Disk data (250 bytes each) being returned, the additional memory usage would be: 10,000 * 250 = 2,500,000 bytes. In sizing the memory requirements for the monitoring server, consider increasing the memory estimate by 50% to allow room for dynamic memory requests.
- See the Jazz for Service Management hardware and software requirements at http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install_prereqs.html.

Additional requirements

This section contains some additional system requirements.

- The best network connection possible is needed between the hub monitoring server and portal server and also between the Tivoli Data Warehouse, Warehouse Proxy Agent, and Summarization and Pruning Agent.

- A video card supporting 64,000 colors and 1024 x 768 resolution is required for the portal client.

Required hardware for System z

The Tivoli Enterprise Monitoring Server can be installed on either a z/OS or Linux operating system running on System z hardware. The Tivoli Enterprise Portal Server is supported on Linux for zSeries, but not z/OS.

The supported product versions for z/OS and Linux for zSeries are listed in Table 11 on page 77.

The following paragraphs summarize the hardware requirements for IBM Tivoli Monitoring components that run on zSeries:

- Tivoli Enterprise Monitoring Server on z/OS

The Tivoli Enterprise Monitoring Server can run natively on any zSeries hardware under z/OS V1.10 or later.

If you enable the self-describing agent feature in your environment, your monitoring servers must have access to a HFS or zFS file system through UNIX System Services (USS). A recommendation is to allocate 25 to 50 MB for this file system. The amount of storage required for the packages of a self-describing agent depends on the monitoring server type:

- Hub monitoring server = 3.25 MB
- Remote monitoring server = 2 MB

This average number takes into account the following factors:

- The average amount of storage needed by the monitoring server self-describing agent process to backup the existing monitoring server product files in the \$TEMS_MANIFEST_PATH/SDMBACKUP directory.
- The hub monitoring server self-describing agent process stores all existing self-describing agent package files including a self-describing agent's TEP(S) JAR files, which are the largest JAR files, often in the 1 to 2 MB range.
- The remote monitoring server only has to store a self-describing agent's monitoring server JAR files, which are typically around 50K.

The amount of USS disk space required is also influenced by the following additional factors:

- The amount of space fluctuates depending on how many self-describing agents register with the z/OS monitoring server, and whether multiple self-describing agent installations are in progress.
- The JAR files are extracted into individual files in subdirectories under the monitoring server USS home directory, and those individual files are copied to &rhilev.&rte.RKANDATV and then automatically deleted. You must allocate enough space for both the JAR files and the extracted contents of those JAR files. The recommended USS disk space must be sufficient to handle these spikes in disk usage.

- Tivoli Enterprise Monitoring Server or Tivoli Enterprise Portal Server on Linux for zSeries

Any zSeries hardware provides adequate processing capability for a monitoring server or portal server running under Linux for zSeries. The Linux operating system must be configured to utilize a 64-bit kernel. The memory and disk requirements for the Tivoli Monitoring components on Linux for zSeries are similar to those for Linux running on distributed systems; these are documented in Table 18 on page 85.

Required software

The following table lists the software required for IBM Tivoli Monitoring.

Table 19. Required software for IBM Tivoli Monitoring

Product	Supported version	Component where the software is required				
		Monitoring server	Portal server	Portal desktop client	Portal browser client	Monitoring agent
Distributed systems: IBM Runtime Environment for Java	JRE V1.6		³		X ¹	
Distributed systems: IBM Runtime Environment for Java	JRE V1.7	X ³	X ³	X ³	X ¹	
Distributed systems: Oracle Java SE Runtime Environment	JRE V1.6.xx, V1.7.xx			X	X	
z/OS systems: IBM 31-bit or 64-bit SDK for z/OS, Java Technology Edition	IBM Java SDK for z/OS at V6 or later	X ⁴				
z/OS systems: z/OS UNIX System Services		X ^{4 5}				
z/OS systems: Integrated Cryptographic Service Facility		X ⁶				X ⁷
For Linux computers: a Korn shell interpreter	ksh package provided by the Linux distribution or pdksh-5.2.14	X	X	X		X ²
AIX only: xlc Runtime Environment	Component xlc.aix61.rte level 11.1 is required for monitoring server and portal server. For monitoring agent, level 9 or higher is required.	X	X			X
Database	<p>A supported RDBMS is required for the Tivoli Enterprise Portal Server and the Tivoli Data Warehouse. Supported database platforms for the portal server are listed in Table 14 on page 82. Supported database platforms for the Tivoli Data Warehouse are listed in Table 15 on page 83.</p> <p>Each database requires a driver. For detailed information, see Chapter 19, "Tivoli Data Warehouse solutions," on page 509 and subsequent chapters about the Tivoli Data Warehouse.</p>					

Table 19. Required software for IBM Tivoli Monitoring (continued)

Product	Supported version	Component where the software is required				
		Monitoring server	Portal server	Portal desktop client	Portal browser client	Monitoring agent
Note:						
<ol style="list-style-type: none">1. If the JRE is not installed on the computer on which the browser is launched, you are prompted to download and install it from the portal server. The Windows user account must have local administrator authority to download and install the JRE. The portal browser clients can use either the 32-bit or 64-bit JRE at version 1.6 or 1.7. The default JRE available from the portal server is dependent on whether the portal server was installed or upgraded. On a pristine install of the portal server, IBM JRE 1.7 is the default JRE for portal browser clients. However, if the portal server was upgraded, IBM JRE 1.6 remains the default JRE.2. The Korn shell (any version) is also required when installing the monitoring agent on AIX, HP-UX, or Solaris systems.3. IBM JRE 1.7 is installed with the monitoring server, portal server, and portal desktop client components on distributed systems. IBM JRE 1.6 is also installed on the portal server for use by portal browser clients.4. If the self-describing agent function is enabled, you must install this software on the z/OS systems where your monitoring servers are installed.5. A file system (either HFS or zFS) and user security must be configured for UNIX System Services on the z/OS systems where monitoring servers are installed in order to use the self-describing agent function.6. See “Global Security Toolkit” on page 142 for configuration steps that are required if Integrated Cryptographic Service Facility (ICSF) is not installed and enabled on a z/OS system with a monitoring server. Cryptographic coprocessor cards are not required to use ICSF.7. If the Integrated Cryptographic Service Facility (ICSF) is not installed and enabled on a z/OS system with a monitoring agent using the Tivoli Enterprise Monitoring Agent framework at V6.3 or later, the monitoring agent cannot validate security tokens included with Take Action requests but the agent will still execute the Take Action request. (The security tokens are used to verify that the request originated from a monitoring server or portal server in your environment and to generate an audit record identifying the user who initiated the request.) Cryptographic coprocessor cards are not required to use ICSF.						

Supported browser versions

The following table shows the supported browser versions for the Tivoli Enterprise Portal browser client, Infrastructure Management Dashboards for Servers application in the Dashboard Application Services Hub, and the Tivoli Monitoring Installation Launchpad.

Table 20. Supported browser versions

Browser version	Supported version	Tivoli Enterprise Portal browser client	Infrastructure Management Dashboards for Servers application	Tivoli Monitoring Installation Launchpad
Microsoft Internet Explorer	V8.0 with all critical Microsoft updates applied	X	X	X
Microsoft Internet Explorer	V9.0 with all critical Microsoft updates applied	X	X	X
Microsoft Internet Explorer	V10.0 with all critical Microsoft updates applied	X		X
Mozilla Firefox	Extended Support Release 10	X		X

Table 20. Supported browser versions (continued)

Browser version	Supported version	Tivoli Enterprise Portal browser client	Infrastructure Management Dashboards for Servers application	Tivoli Monitoring Installation Launchpad
Mozilla Firefox	Extended Support Release 17	X	X	X

Required software for event integration with Netcool/OMNIBus

This section contains the required software for event integration with Netcool/OMNIBus.

The following products must be installed and configured before you install event synchronization and configure event forwarding for Netcool/OMNIBus:

- IBM Tivoli Netcool/OMNIBus Probe for Tivoli EIF version 10 or later and the non-native probe version 12 or later
- Netcool/OMNIBus 7.2.1 FP10 or later fix pack
- Netcool/OMNIBus 7.3.0 FP5 or later fix pack
- Netcool/OMNIBus 7.3.1 or later fix pack
- Netcool/OMNIBus 7.4 or later fix pack

Note:

1. If you are installing IBM Tivoli Monitoring event synchronization on a Linux operating system, you must ensure that the libXp and libXtst 32-bit shared libraries are installed on the computer system before installing event synchronization.
2. If you want to configure SSL between the probe and monitoring agents, version 12.0 or later of the Netcool/OMNIBus Probe for Tivoli EIF is required.
3. If you are using the virtualization rules and triggers and predictive rules and triggers provided with Netcool/OMNIBus, you must use Netcool/OMNIBus Version 7.3.0 FP 6 or later, or Netcool/OMNIBus Version 7.3.1 FP2 or later, as those fix packs have the required updates to the rules files and triggers for integration with IBM Tivoli Monitoring.

Required software and memory requirements for a dashboard environment

If you are installing a dashboard environment with monitoring dashboard applications such as the IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM, or creating custom monitoring dashboards, and optionally using the Tivoli Authorization Policy Server, install the following Jazz for Service Management version 1.1 (or later) components:

- IBM WebSphere Application Server V8.5.0.1

Note:

1. If you already have IBM WebSphere Application Server V8.0.0.2 or a later fix pack, or WebSphere Application Server V8.5.0.1 or a later fix pack in your environment, you can use it instead of installing WebSphere Application Server V8.5.0.1 from the Jazz for Service Management media. See Installing Jazz for Service Management for more details on installing into an existing WebSphere Application Server (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html).

2. The dashboard environment must be installed on a Dashboard Application Services Hub that was installed on a WebSphere Application Server Full Profile. A Dashboard Application Services Hub that is installed on a Liberty Profile is not supported by IBM Tivoli Monitoring.

See the WebSphere Application Server detailed system requirements at: <http://www-01.ibm.com/support/docview.wss?rs=180&uid=swg27006921>.

- Jazz for Service Management extension for IBM WebSphere
- Dashboard Application Services Hub V3.1

See Installing Jazz for Service Management (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html) for the following information:

- Hardware, operating system, and software requirements for Dashboard Application Services Hub.
- Installation technologies that are supported for the Jazz for Service Management components. Jazz for Service Management provides a launchpad application to guide you through a typical or custom installation procedure and also supports installation of Dashboard Application Services Hub from an IBM Installation Manager local or network repository.

IBM Tivoli Monitoring includes the Jazz for Service Management media.

After installing and configuring the Jazz for Service Management prerequisite components that are listed above, you can use the IBM Installation Manager Version 1.6.1 or later or the IBM Tivoli Monitoring installation launchpad to install the following IBM Tivoli Monitoring components into the Dashboard Application Services Hub:

- IBM Infrastructure Management Dashboards for Servers application
- Tivoli Authorization Policy Server

The tivcmd Command Line Interface (CLI) for Authorization Policy component should also be installed on the computer systems that are used by the administrators who create and work with authorization policies. The tivcmd CLI does not require any Jazz for Service Management components to coexist on the same system. You must use IBM Installation Manager version 1.6.1 or later or IBM Tivoli Monitoring installation launchpad to install the tivcmd CLI.

See “Using IBM Installation Manager to install or update IBM Tivoli Monitoring components for a dashboard environment” on page 299 for more details on the installation options for installing the Infrastructure Management Dashboards for Servers application, the authorization policy server, and the tivcmd CLI. Detailed instructions for installing each of these components are available here:

- “Installing and configuring the IBM Infrastructure Management Dashboards for Servers” on page 306
- “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318

Note:

1. A minimum of 200-300 MB of free space is required on the system where you are installing the Tivoli Authorization Policy Server.
2. On Linux systems, you must install the GTK (GNU Image Manipulation Program Toolkit) support packages before you use the IBM Tivoli Monitoring installation launchpad or the IBM Installation Manager graphical user interface. GTK toolkit is a multi-platform toolkit for creating graphical user interfaces. The package get utilities for your Linux distribution can be used to install the latest versions of the GTK libraries and the GTK library dependencies.
3. Both the Jazz for Service Management components and the IBM Tivoli Monitoring components for a dashboard environment must be installed in the same user mode; install both as a root user or install both as a non-root user. Information cannot be exchanged between these components when there is a combination of root and non-root installations.

4. On Linux or AIX systems, you must install the Korn shell interpreter (ksh) before you install the Tivoli Authorization Policy Server or the tivcmd Command Line Interface for Authorization Policy. Either use the ksh package provided by the Linux distribution or the public domain Korn shell package pdksh-5.2.14.
5. If you are using the IBM Tivoli Monitoring installation launchpad to perform the installation or upgrade of components, a Web browser is required on the target computer. For more information, see Table 20 on page 90.

In addition to installing the Jazz for Service Management components and IBM Tivoli Monitoring components, you must also perform the following configuration tasks:

1. You should also install an LDAP server and configure the WebSphere Application Server for Dashboard Application Services Hub and your Tivoli Enterprise Portal Server to use a federated user repository and single sign-on.

Note: An LDAP server and single sign-on are not required if all dashboard operators can use the same permissions for display monitoring resources in Dashboard Application Services Hub and you do not need single sign-on support for launching to the Tivoli Enterprise Portal.

2. Enable the dashboard data provider when you configure the Tivoli Enterprise Portal Server. From Dashboard Application Services Hub, configure a connection to the dashboard data provider.
3. If you are using the Authorization Policy Server, you must also enable authorization policies during portal server configuration. However, the recommended approach is to set up your dashboard environment with authorization policies disabled. After you have a dashboard user who can view monitoring data by using the IBM Infrastructure Management Dashboards for Servers application, or custom dashboards, and your administrators have created authorization policies for your dashboard operators, then reconfigure the portal server to enable authorization policies.

For more information about the configuration tasks to perform after the installation tasks are complete, see the *IBM Tivoli Monitoring Administrator's Guide* section on *Preparing your dashboard environment*. See also "Supported browser versions" on page 90 for the list of browsers that are supported by IBM Infrastructure Management Dashboard for Servers.

If you have IBM SmartCloud Monitoring version 7.2 or later or IBM Tivoli Monitoring for Virtual Environments version 7.2 or later, you can install their dashboard monitoring applications (IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM) into Dashboard Application Services Hub. You can also install any other monitoring dashboard applications that use the IBM Tivoli Monitoring V6.3 or later dashboard data provider.

Required software and memory requirements for a reporting environment that uses Tivoli Common Reporting

If you are planning to use Tivoli Common Reporting and report packages that are provided with monitoring agents or custom reports for monitoring agents, first install and configure the Tivoli Data Warehouse. For more information, see Chapter 19, "Tivoli Data Warehouse solutions," on page 509.

You can install the reports for IBM Tivoli Monitoring Version 6.3 operating system agents with Tivoli Common Reporting 3.1 or 2.1.1. IBM Tivoli Monitoring V6.3 includes Tivoli Common Reporting 3.1, so you must use version 3.1 when to set up a new reporting environment. Tivoli Common Reporting 3.1 is a Jazz for Service Management component. To use this version of Tivoli Common Reporting, install and configure the following Jazz for Service Management components if they are not already in your environment:

- IBM WebSphere Application Server V8.5.0.1

Note: If you already have IBM WebSphere Application Server V8.0.0.2 or a later fix pack or WebSphere Application Server V8.5.0.1 or a later fix pack in your environment, you can use it instead of installing

WebSphere Application Server V8.5.0.1 from the Jazz for Service Management media. See the Jazz for Service Management Installation Guide for more details on installing into an existing WebSphere Application Server.

- Jazz for Service Management extension for IBM WebSphere
- Dashboard Application Services Hub V3.1
- Tivoli Common Reporting V3.1

For information about the WebSphere Application Server system requirements, see <http://www-01.ibm.com/support/docview.wss?rs=180&uid=swg27006921>.

See Installing Jazz for Service Management (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html) for the following information:

- System requirements for Dashboard Application Services Hub and Tivoli Common Reporting.
- Installation technologies that are supported for the Jazz for Service Management components. Jazz for Service Management provides a launchpad application to guide you through a typical or custom installation procedure. Jazz for Service Management also supports installation of Dashboard Application Services Hub from an IBM Installation Manager local or network repository. You must use the Tivoli Common Reporting installation program to install Tivoli Common Reporting.

After installing and configuring Tivoli Common Reporting and its prerequisites, go to Chapter 28, “Using Tivoli Common Reporting and monitoring agent reports,” on page 693 for more information.

Required software and memory requirements for using the OSLC Performance Monitoring service provider

If you plan to use the IBM Tivoli Monitoring OSLC Performance Monitoring service provider to integrate with other products by using OSLC linked data interfaces, you must install the following Jazz for Service Management version 1.1 prerequisite components if they are not already in your environment:

- IBM WebSphere Application Server V8.5.0.1

Note: If you already have IBM WebSphere Application Server V8.0.0.2 or a later fix pack or WebSphere Application Server V8.5.0.1 or a later fix pack in your environment, you can use it instead of installing WebSphere Application Server V8.5.0.1 from the Jazz for Service Management media. See the Jazz for Service Management Installation Guide for more details on installing into an existing WebSphere Application Server.

- Jazz for Service Management extension for IBM WebSphere
- DB2 database server for the Registry Services database if you do not already have a database server in your environment
- Registry Services V1.1
- Security Services V1.1 (optional)

Note:

1. IBM Tivoli Monitoring includes the Jazz for Service Management version 1.1 media.
2. If you want the Performance Monitoring service provider to authenticate HTTP requests from OSLC clients, you should install and configure Security Services. Security Services enable non WebSphere based applications such as the Performance Monitoring service provider to participate in LTPA based single sign-on. See “Single sign-on capability” on page 141 for more details on using Security Services and single sign-on with the Performance Monitoring service provider.
3. Registry Services and Security Services must be installed into the same WebSphere Application Server.
4. Registry Services support two types of authentication:

- a. Basic authentication which requires users and client applications to provide a user ID and password to Registry Services.
- b. HTTP client certificate authentication which requires the use of HTTPS and also requires users and client applications to provide a client certificate to Registry Services.

Because the Performance Monitoring service provider cannot be configured to provide an HTTP client certificate to Registry Services, you must configure Registry Services to use basic authentication when authentication is enabled. In this case, you must also configure Tivoli Enterprise Monitoring Automation Server with a user ID and password that you can use to authenticate with Registry Services.

For information about the WebSphere Application Server system requirements, see <http://www-01.ibm.com/support/docview.wss?rs=180&uid=swg27006921>.

See Installing Jazz for Service Management (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html) for the following information:

- System requirements for Registry Services and Security Services.
- Installation technologies that are supported for the Jazz for Service Management components. Jazz for Service Management provides a launchpad application to guide you through a typical or custom installation procedure and also supports installation from an IBM Installation Manager file or network repository.

After you install and configure Registry Services and optionally Security Services, install and configure the Tivoli Enterprise Monitoring Automation Server. As part of the configuration, you configure connection information for Registry Services. For more information, see “Installing and configuring Tivoli Enterprise Monitoring Automation Server” on page 219.

Required software and memory requirements for non-linear trending in Tivoli Performance Analyzer

This section contains the software and memory requirements for non-linear trending in Tivoli Performance Analyzer.

The following are prerequisite criteria for using the non-linear trending feature in Performance Analyzer:

- You must install SPSS® Forecast Server V20.
- A minimum of 100 MB of free space is required in the ITPA agent home directory for the output data files generated by SPSS.

Chapter 4. Deployment phase

Proper planning ensures that the deployment phase goes smoothly.

This chapter reviews the steps necessary to deploy and configure your Tivoli Monitoring environment. During the initial installation, use a phased approach. Perform these steps in development and test environments before applying them in the production environment.

Pre-installation checklist

Use the following checklist for your pre-installation.

- Verify that your hardware meets the requirements for Tivoli Monitoring core components.
- Verify that the correct media is downloaded (not the upgrade install media).
- Verify that you have the correct media for your hardware architecture. Some operating systems support both 32-bit and 64-bit kernels. To check which kernel version your system is running, use the following commands:

Table 21. Commands for determining your system's kernel version

System	Command
AIX	bootinfo -K get 64 or bootinfo -K get 32
HP	getconf KERNEL_BITS
Linux	uname -r
Solaris	isainfo -b

- Verify prerequisites are met for the databases.
- Confirm you have the appropriate authority on the system to be installed.
- Confirm the network team is contacted for required ports to be opened between the portal server, monitoring server, monitoring agents, and Warehouse Proxy Agents.
- Ensure the team responsible to deploy and implement Tivoli Monitoring solution is equipped with sufficient skills to complete the tasks successfully.

Give special consideration to Windows installations where you are using terminal services or some other *remote control* software to access hosts. In a perfect world, install this software while you are physically sitting at and logged in to the system console. As this might be impractical, you need to have LOCAL ADMINISTRATOR authority (domain administrator authority itself does not suffice).

If using something like terminal services, you need to be a member of the local user security policy called *Allow Logon Locally* as well as *Allow Logon through Terminal Services*. In some cases, you might need to work with the Windows engineers to help you maneuver through local and AD security issues. Security problems can initially manifest themselves as subtle inconsistencies but ultimately can inhibit a successful installation.

Installing the infrastructure components

Before installing any Tivoli Monitoring agents, you must first install the Tivoli Monitoring infrastructure components, which include the hub monitoring server, portal server, remote monitoring server, Warehouse Proxy Agents, and Summarization and Pruning Agent.

For the first phase of your deployment, deploy a relatively small environment, which provides you with the opportunity to test your production environment and ensure that it is running smoothly. Deploy the

hub monitoring server and only two remote monitoring servers initially. The remote monitoring servers can successfully monitor up to 1500 agents, so using two remote monitoring servers during your initial deployment is not a problem. Firewall considerations might necessitate additional remote monitoring servers.

If you plan to use clustering to achieve a high availability configuration, configure the cluster before connecting any agents. Otherwise, it is necessary to reconfigure the agents to connect to the cluster rather than connecting to one of the nodes in the cluster. For information on setting up the hub monitoring server and portal server in an OS, see the *IBM Tivoli Monitoring High Availability Guide for Distributed Systems*.

Because the installation of application agent support on the monitoring server and portal server requires them to be recycled, install the application support on the monitoring server, portal server, and portal client for all agents that you expect to deploy in the next six to nine months. This likely includes the IBM Tivoli Monitoring for Databases, IBM Tivoli Monitoring for Messaging and Collaboration, ITCAM for SOA and any other products you plan to implement.

With the self-describing agent automated application support feature, an automated installation process is available which is internal to the IBM Tivoli Monitoring agent, Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, and Tivoli Enterprise Portal client. The self-describing agent capability must be correctly configured for a product before deployment. For more information, see "Enabling self-describing agent capability at the hub monitoring server" on page 204 and the *IBM Tivoli Monitoring Administrator's Guide*.

To take advantage of the self-described agent capability, you must have the Tivoli Management Services at Version 6.2.3 (or higher).

During the initial installation phase, install the Warehouse Proxy Agent and Summarization and Pruning Agent, but do not begin using them until after you have successfully deployed your first 50 agents. Installing in this way gives you an opportunity to assess the health of your environment before adding to the complexity of your environment.

Configuration checklist

Use the following checklist for your configuration.

- Install all Tivoli Monitoring core components (portal server, monitoring server, Warehouse Proxy Agent, and Summarization and Pruning Agent) so there can be a section for each component.
- Verify the correct protocols are selected. If SPIPE is chosen, make sure the encryption key string used is the same across the Tivoli Monitoring enterprise environment.
- Verify the correct configurations are performed regarding data warehousing.

After installing the hub and remote monitoring server, ensure that you do not attempt to start a second kdsmain instance, which can corrupt your environment. Modify the monitoring server startup *CandleServer* script so that it looks as follows:

```
#
#Local change to check for another running kdsmain
#
if [ "$action" = "start" ]
then
    if ps -ef | grep -v grep | grep kdsmain
    then
        echo "There is a KDSMAIN running already"
        exit
    fi
fi
```

Some users run multi-hub monitoring servers on a single server for the development and test environments using different ports. If that is the case, then the previous script cannot work because `kdsmain` is running for the other hub. If that is the case, use extreme care to ensure that you do not accidentally start a second `kdsmain` for a given hub.

Do *not* enable warehousing at this time. Wait until all the agents for phase one have been installed and started, and the situations have been distributed.

Disable all the default situations by unassigning the managed system group and any agents present in the Assigned check box.

Create the managed system groups before creating the situations.

Distribute all newly created situations with naming conventions to the customized managed system group and not to `*NT_SYSTEM`, `*ALL_UNIX`. You must customize your situation thresholds before forwarding your events to Tivoli Enterprise Console or Netcool/OMNIBus, which ensures that you do not cause any event storms.

You can enable event forwarding at this time. Install your first 50 agents.

Using the installation method of your choice, install several OS monitoring agents. These can be installed locally on the server or through the remote deployment mechanism. Detailed steps on remote deployment are included in the following sections.

Note: Because the remote deployment of application agents depends on having an OS agent running on the server, always deploy the OS agents first.

Customizing your environment

You can customizing situations, workspaces, and queries.

Before progressing any further, you must perform customization to your environment. The following sections describe this customization.

Customizing situations, workspaces, and queries

Perform some customizations of your situations to ensure that this first set of agents do not generate event storms. Many users choose to disable the default situations and then create their own situations. Careful thought must be put into the naming conventions of your situations. Many users include the following elements in their situation names.

- OS type
- Agent name or type, or both
- Business unit
- Physical location
- Severity

When choosing a name, keep in mind that the situations are sorted alphabetically in the Situation Editor. A typical situation might look like:

- `East_UNIX_High_CPU_Crit`

For more information on disabling the default situations and performing bulk work on situations see the *IBM Tivoli Monitoring Command Reference*.

Choose similar naming conventions for any custom queries, managed system groups, and workspaces. Choose the same criteria such as physical location, business unit, agent type, and so on that you used for situations.

Another important consideration when creating situations is the Display Item, which by default, is not enabled. If you want to generate a unique Tivoli Enterprise Console Event for each item that triggered a situation or want to run a Take Action command against each item that triggered the situation, then you want to select the Display Item and choose the appropriate attribute.

Changing the default monitoring server configuration settings

There are a few configuration settings that you might want to change from the default values.

Consider changing the following settings to ensure a smooth running environment even as your environment scales.

- **KDCFC_RXLIMIT**

- This is the buffer used for return queries, specified in KB. The default value is 2048 KB (equivalent to 2 MB). A value of 8192 KB (equivalent to 8 MB) seems to work well for most users.
- Recommendation: 8192 KB

- **DEPOTHOME**

- The location of the depot. The default location is

Windows %CANDLE_HOME%\CMS\Depot

Linux **UNIX** \$CANDLEHOME/tables/hub_tems_name/depot

- Relocating the depot directory enables you to backup your Tivoli Monitoring environment without having to backup the very large depot directory. In addition, relocating the depot directory ensures that the depot will not fill up the filesystem where Tivoli Monitoring is running.

The target directories listed below are examples:

If CANDLE_HOME on **Windows** is located at C:\IBM\ITM, then relocate the depot to D:\ITM\depot

If CANDLE_HOME on **Linux** **UNIX** is located at /opt/IBM/ITM, then relocate the depot to /data/ITM/depot

- Bind a specific IP address

KDEB_INTERFACELIST=192.100.100.100

Note: Use this option only if the monitoring server and portal server are on separate servers. If they are on the same computer, there are going to be problems due to the multiplexed port 1920: The tacmd command will not be able to find the monitoring server, and portal server clients will not be able to find the portal server.

- Bind a specific host name

KDEB_INTERFACELIST=caps001

- Bind the first IPV4 address associated with the current host name to be the default interface.

KDEB_INTERFACELIST=!*

There are many other optional parameters documented in the *IBM Tivoli Monitoring Administrator's Guide*. Review those parameters and determine whether any are required for your environment. See "Tivoli Enterprise Monitoring Server" on page 470 for information on monitoring server parameters for performance tuning.

Installing your first 50 agents

You can install agents locally on the server or through the remote deployment mechanism

Using the installation method of your choice, install several OS monitoring agents. These agents can be installed locally on the server or through the remote deployment mechanism. Because the remote deployment of application agents depends on having an OS agent running on the server, always deploy the OS agents first.

Remote deployment of OS agents

Tivoli Monitoring provides default *CreateNode* and *AddSystem* flags using the *tacmd* command to deploy OS agents as well as application agents remotely from a central location. It is very important that the depots be created on each monitoring server (hub and remote) and that you make sure the level of code, bundles, and packages in those installation depots is consistent. You can also use the shared depot. If a shared directory is accessible to all monitoring servers, it can be mounted across the monitoring server environment. This reduces the maintenance workload, since you have only one directory to maintain in a shared location, rather than maintaining depot directories on each monitoring server.

Postinstallation checklist

A postinstallation checklist to ensure you have completed the important deployment items.

Use this postinstallation checklist to ensure the following items have been completed:

- Monitoring server (situations created)
- Portal server (check all aspects of portal server functionality such as workspaces)
- Perform a complete backup of all Tivoli Monitoring components
- If the self-describing agent feature is enabled, use the command `tacmd listappinstallrecs -t <pc>` to verify that self-describing agent support installation was successful after the agent connected to the infrastructure.

Configuring your warehouse agents

Now that you have your first 50 agents up and running smoothly, it is time to configure your Warehouse Proxy Agent and Summarization and Pruning Agent.

There are the two major components to configuring your warehousing.

- First, configure the two agents and specify all of the steps necessary for the agents to communicate with the database server. These steps are done through the configuration panels that appear when you reconfigure the agents. Follow the steps in the install guide when performing this configuration. Choose the summarization and shift options that you chose during the planning phase.
- The second aspect to warehousing is to decide which attribute groups you are going to collect and at what intervals. Always use the Warehouse load projection spreadsheet before enabling additional historical collection. Thus, ensuring that you do not overload your warehouse environment. Start slowly and incrementally add attribute groups to your historical collection. This allows you to confirm that everything is working properly before you add more data to your warehouse.

Enable one attribute group at a time and verify that the data is being collected, summarized and pruned properly. At this point, you can also see how many rows are getting written per data collection interval for the attribute group, by examining the entries in the WAREHOUSELOG table. The number rows written per interval is an important input parameter for the Warehouse load projection spreadsheet. If the WAREHOUSELOG is disabled, you can look at the self monitoring Warehouse Proxy Agent workspace. The workspace displays the top ten nodes with the greatest number of exports since the Warehouse Proxy Agent started, and the ten most recent errors in the last 24 hours.

At this point, install the warehouse monitoring solution that is available at the IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>) by searching for “Data Warehouse DB activity” or navigation code “1TW10TM1X”.

In addition, you must create a critical situation that monitors both your Warehouse Proxy Agent and Summarization and Pruning Agent to ensure that they are running. If they are not running, the situation must run a Take Action to automatically restart those agents.

When configuring and starting your Summarization and Pruning Agent, do not forget to set the **KSY_MAX_WORKER_THREADS** to the appropriate number for your environment. See “Locating and sizing the Summarization and Pruning Agent” on page 39 for recommendations on the number of worker threads.

If you have configured some or all of your monitoring agents to run autonomously, you might want to configure the Warehouse Proxy and Summarization and Pruning agents to run autonomously as well. See “Running the warehouse agents autonomously” on page 658 for more information.

Installing additional agents

After the first 50 agents are deployed, deploy additional agents in groups of approximately 200, which allows you to verify all of the agents before progressing to the next step.

At this point in your deployment, keep things simple by deploying a single hub monitoring server, a portal server, the warehouse components, and a couple of remote monitoring servers. If warehousing is critical, it can be enabled at this point. The environment must be evaluated thoroughly before moving beyond the initial 50 agents.

As you scale up your environment and become comfortable with the deployment and validation of agents, you can increase the number of agents being deployed to 400 or 500 at a time.

Chapter 5. Tivoli Monitoring Startup Center

The IBM Tivoli Monitoring Startup Center is a graphical user interface tool that guides you through the setup of a new IBM Tivoli Monitoring environment that includes the following components.

- Tivoli Enterprise Monitoring Server
- Tivoli Enterprise Portal Server
- Tivoli Enterprise Portal browser client
- Operating system agent
- Warehouse Proxy Agent
- Remote Tivoli Enterprise Monitoring Server (optional component)
- Tivoli Enterprise Portal desktop client (optional component)

The Startup Center cannot be used to deploy and configure the following monitoring components:

- Tivoli Enterprise Monitoring Automation Server and its Performance Monitoring service provider
- A dashboard environment that includes Dashboard Application Services Hub, monitoring dashboard applications such as IBM Infrastructure Management Dashboards for Servers, Authorization Policy Server, and the tivcmd Command Line Interface for Authorization Policy.

The Startup Center is not intended for upgrading existing IBM Tivoli Monitoring components. You use topology diagrams to configure and deploy an initial base IBM Tivoli Monitoring environment, which can be expanded later. The Startup Center reduces complexity, increases transparency, and simplifies your IBM Tivoli Monitoring deployment.

The Startup Center can run on both Windows and Linux Intel x86-32 systems. On Linux systems you must have installed the GTK (GNU Image Manipulation Program Toolkit) support packages to run the tool. GTK toolkit is a multi-platform toolkit for creating graphical user interfaces. The IBM Tivoli Monitoring product image location that is used in the deployment must be accessible.

In the event of a failed installation, the Startup Center log files can help to diagnose the problem. The log files can be found in the following locations:

- **Windows** %USERPROFILE%\STARTUP\logs\
- **Linux** **UNIX** \${HOME}/.STARTUP/logs

Startup Center platform support

This section contains platform support information for the Tivoli Monitoring Startup Center.

Table 22 on page 104 and Table 23 on page 104 include a list of platform support for launching the Startup Center or installing it locally on Windows and Linux Intel x86-32 operating systems. For remote installation on distributed platforms and the supported monitoring components, see “Supported operating systems” on page 75. For more information on platforms no longer supported in V6.3, see “Platforms no longer supported for IBM Tivoli Monitoring V6.3” on page 145.

Note:

1. On **Windows** the Startup Center installs 32 bit components to 32 bit systems and 64 bit components to 64 bit systems.
2. The Startup Center is not supported on **z/OS** installations.

Table 22. Startup Center: Supported Windows operating systems

Operating system	Launch Startup Center or Install locally
Windows 7 on x86-32	X
Windows 7 on x86-64	X

Supported Linux operating systems

Table 23. Startup Center: Supported Linux operating systems

Operating system	Launch Startup Center or Install locally
RedHat Enterprise Linux 5.0 Intel x86-32 with GTK support	X
SuSE Linux Enterprise Server 11 Intel x86-32 with GTK support	X

Startup Center prerequisite requirements

This section contains prerequisite information for the Tivoli Monitoring Startup Center.

For a distributed installation, you must have DB2 already installed on the target machine where you want to install the Tivoli Enterprise Portal Server and the Warehouse Proxy Agent.

Remote Execution and Access (RXA) is an IBM developer toolkit that provides classes and methods to create an application that can establish a connection with a remote computer. The Startup Center relies on RXA to establish a connection with a remote computer. Therefore you must enable the target machines for RXA.

Enabling the RXA toolkit on UNIX and Linux systems

- You must ensure that the SSH protocol is installed and enabled on any target that you want to access by using SSH protocol. OpenSSH 3.71 (or higher) contains security enhancements that are not available in earlier releases.
- RXA cannot establish connections with any UNIX target that has all remote access protocols (rsh, rexec, or ssh) disabled.
- For RXA to communicate with Linux and other SSH targets by using password authentication, you must edit the `/etc/ssh/sshd_config` file on target computers and set the parameter `PasswordAuthentication` to `yes`. The default setting is `no`. After changing this setting, stop and restart the SSH daemon by using the following commands:

```
/etc/init.d/sshd stop
/etc/init.d/sshd start
```

On AIX systems, you can stop and restart the SSH daemon by using the following commands:

```
/etc/rc.d/rc2.d/Ssshd start
/etc/rc.d/rc2.d/Ksshd stop
```

Enabling RXA on Windows systems:

- Interprocess communications share (IPC\$) is required to enable RXA on Windows systems. You can use the `net share` command to check the status of IPC\$:


```
C:\Documents and Settings\Administrator>net share
```

If IPC\$ is not set correctly, you can use the command `net share C$=C:` to set the IPC\$.

- On Windows Server 2008 and Windows Vista, you might need to disable User Account Control if your account is not a domain user account.

- The account used to install IBM Tivoli Monitoring components on target machines must be in the Administrator group.

For more information, see “Remote Execution and Access” on page 115.

Deployment procedure

This section contains information on launching the Tivoli Monitoring Startup Center.

Procedure

1. First select the appropriate Tivoli Monitoring Startup Center CD for your platform:
 - CD for **Windows** IBM Tivoli Monitoring Startup Center for Windows 32-bit, Multilingual.
 - CD for **Linux** IBM Tivoli Monitoring Startup Center for Linux 32-bit, Multilingual.
2. If you are downloading the Startup Center from the Passport Advantage® Web site, you must first unpack the Startup Center packages. To launch the Startup Center:
 - **Windows** use: launchStartup.bat.
 - **Linux** use: launchStartup.sh.

Results

You have two types of installation to choose from: “Local installation” or “Distributed installation” on page 108.

Note: All screens illustrated in this chapter are for **Windows** systems.

Local installation

A local installation is the simplest form of an IBM Tivoli Monitoring installation and is best suited to small environments and evaluations.

About this task

With a local installation, a predefined set of components are installed on the local machine. You can install the following components on single systems:

- Tivoli Enterprise Monitoring Server
- Tivoli Enterprise Portal Server
- Tivoli Enterprise Portal desktop client
- Tivoli Enterprise Portal browser client
- Operating system agent

Complete the following steps to perform a local installation using the Startup Center:

Procedure

1. Launch the Startup Center.
2. On the Welcome page, click **Next**.



Figure 11. Tivoli Monitoring Startup Center Welcome page

3. Click **Accept** to accept the license agreement.
4. Select the **Local** radio button and click **Next**.

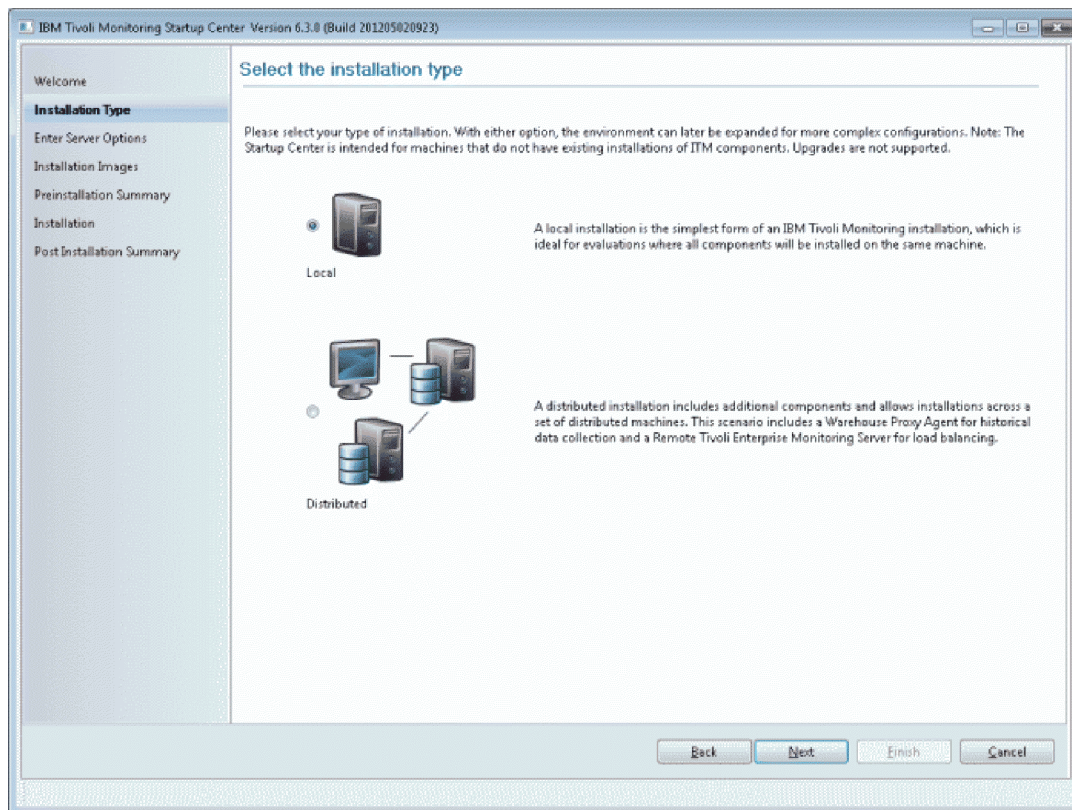


Figure 12. Tivoli Monitoring Startup Center Installation Type window: Select the installation type

5. Enter the installation options for the components in the topology diagram and click **Next**.

Note:

- a. Running multiple installations on a Windows machine specifying different directories is not supported.
- b. On Linux systems, you are not required to enter a password for the embedded Derby database.

In the event of a failed installation, the Startup Center log files can help to diagnose the problem. The log files can be found in the following locations:

- **Windows** %USERPROFILE%\STARTUP\logs\
- **Linux** \${HOME}/.STARTUP/logs

6. You can place all of your installation images on one directory and click **Select Image Repository** to browse to the directory. You can also click the ellipsis button at the end of the Path field to browse to each installation image individually. The directory path containing installation images cannot contain commas or spaces. Click **Next** to continue.

Note: The installation images must be extracted before you set the image location.

7. The Pre-Installation Summary screen is displayed. Any installation errors are displayed here. Review the summary information in this screen window and click **Back** to take corrective action, or click **Next** to continue.
8. The Deployment window shows the deployment status of each component and a progress indicator. Any installation errors are displayed in the Postinstallation Summary window. Click **Next** when the deployment of each component has completed.

Distributed installation

A distributed installation is recommended for medium and large environments that are monitoring more than 500 systems. A distributed installation offers you the most expansion possibilities as your environment grows. The distributed installation is the ideal configuration if you are looking to set up a robust foundation of IBM Tivoli Monitoring.

About this task

With a distributed installation, a predefined set of components are installed on a set of distributed machines. You can install the following components with a distributed installation:

- Tivoli Enterprise Monitoring Server
- Tivoli Enterprise Portal Server
- Tivoli Enterprise Portal browser client
- Operating system agent
- Warehouse Proxy Agent
- Remote Tivoli Enterprise Monitoring Server (optional component)
- Tivoli Enterprise Portal desktop client (optional component)

Note: The TEPS and WAREHOUS databases are created automatically by the Startup Center tool. The `itmuser` account should have a profile that sets up the DB2 environment correctly before running the Startup Center tool.

Complete the following steps to perform a distributed installation by using the Tivoli Monitoring Startup Center:

Procedure

1. Launch the Startup Center.
2. Click **Next** on the Welcome screen.
3. Click **Accept** to accept the license agreement.
4. Select the **Distributed** radio button and click **Next**.




Figure 13. Tivoli Monitoring Startup Center Installation Type window

5. Your environment can now be scanned to discover machines within the IP range that you specify. You can also click **Add System** to specify a system host name or IP address. If you specify a machine by using the **Add System** option you must ensure that the machine is online. Click **Next** to continue.

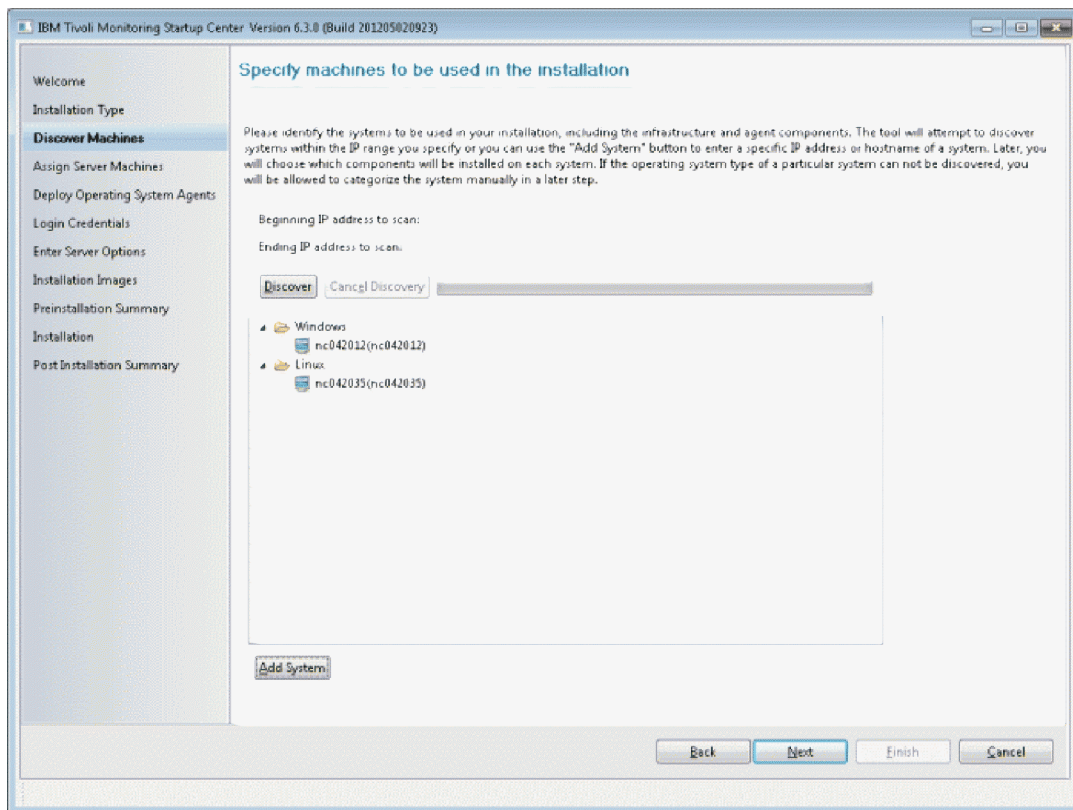


Figure 14. Tivoli Monitoring Startup Center Discover Machines window

Note: For some systems, the Startup Center might not identify the type of operating system. These systems are listed under the **Unknown Operating System** category. For more information, see the *IBM Tivoli Monitoring Troubleshooting Guide*.

6. The Startup Center has discovered the machines in your environment. The components to be installed are represented in a graph. Now you must specify the components to install on a target machine. Select a machine from the list and drag it to the component that you want to install on that machine. You can also press the **Enter** key and select the component from the pop-up dialog box. The Tivoli Enterprise Portal Server and the Warehouse Proxy Agent are assigned to the same machine. If you want to change a machine assignment, select a different machine from the list and drag it to a component. A green check mark indicates a completed machine assignment for the component. Repeat this action for each component in the diagram. You can use the same target machine for more than one component. The Tivoli Enterprise Portal Server, hub Tivoli Enterprise Monitoring Server, and Warehouse Proxy Agent must have a machine assignment before you can proceed. The Tivoli Enterprise Portal Desktop and remote Tivoli Enterprise Monitoring Server are optional components. Click **Next** to continue.

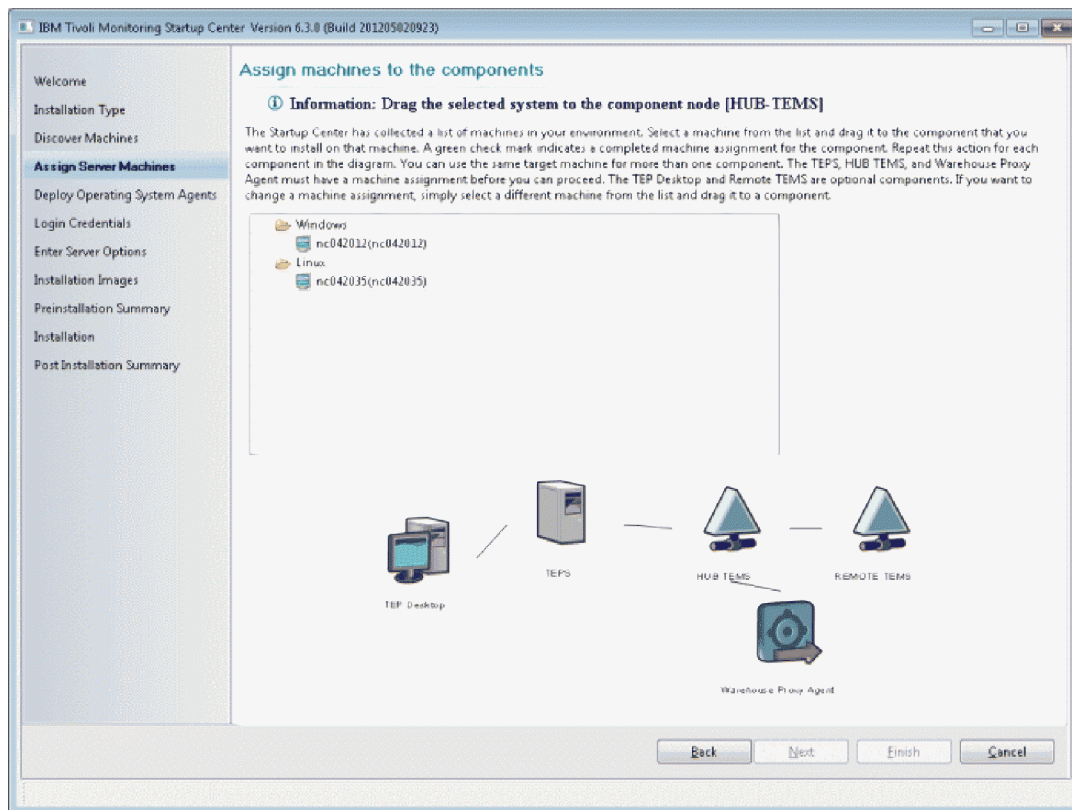


Figure 15. Tivoli Monitoring Startup Center Assign Server Machines window

Note:

- a. If you assign a component to an unknown operating system, you are prompted to specify your operating system.
 - b. By default, when you assign a machine to the Tivoli Enterprise Portal Server, the Startup Center assigns the Warehouse Proxy Agent to the same machine. Likewise, when you assign a machine to the Warehouse Proxy Agent, the Startup Center also assigns the Tivoli Enterprise Portal Server.
 - c. You cannot assign a hub monitoring server and remote monitoring server on the same system.
 - d. You can change a machine assignment but you cannot remove an assignment. If you want to change any of the machine assignments, select another machine from the list and drag it to the component.
7. The Deploy OS Agents screen is displayed. Select the check box of each machine that you want to install an OS Agent on. Enter a path for the target installation directory or accept the default. You must select at least one machine before continuing.

Note: The installation of OS Agents is not supported on Windows Itanium systems. Click **Next** to continue.

8. Now you must provide the machine login credentials that are required to remotely access target systems. Select each machine in the list and enter the System user ID and System password in the fields provided. For information on non-root user credentials, see “Non-root user support” on page 113. You can select multiple machines if they share the same login credentials and enter the information once. All credentials must be validated before you can click **Next**.

Note:

- a. If you encounter a problem at this step, it is probably RXA related. For information about RXA, see “Remote Execution and Access” on page 115.

- b. The **Temporary Directory** must have sufficient space to accommodate the transferred installation image.
 - c. In the Turkish locale, a user ID containing a lowercase "i" causes machine login credentials to fail. You can change to a locale other than Turkish, or else use only capital 'T' letters in user IDs, for example *AdminIstrator* instead of *Administrator*.
9. The Input Server Options screen is displayed. The following components must be configured:
- Tivoli Enterprise Monitoring Server: Enter the installation path for the monitoring server and click **Next**.
 - Remote Tivoli Enterprise Monitoring Server: Enter the installation path for the remote monitoring server and click **Next**.
 - Tivoli Enterprise Portal Server:
 - a. Enter the installation path, password, and database type for the Tivoli Enterprise Portal Server. The database types that are currently supported are embedded Derby and DB2. If you select DB2 as your database type, you must enter the corresponding database information in the fields provided.
 - b. Click **Next**. The database connection is automatically validated.
 - Tivoli Enterprise Portal Server Desktop: Enter the installation path for the Tivoli Enterprise Portal Server Desktop and click **Next**.
 - Warehouse Proxy Agent: Enter a valid installation path along with the corresponding database information and click **Next**.

Note:

At this point, an error might be displayed regarding your connection settings. The RXA tool is used to establish a connection with a remote computer. You must ensure that your target system meets the requirements to establish remote execution and access. For more information see "Remote Execution and Access" on page 115.

If the Startup Center fails to create a warehouse database or database user, see the for more information.

10. You can place all of your installation images on one directory and click **Select Image Repository** to browse to the directory. You can also click the ellipsis button at the end of the Path field to browse to each installation image individually. The directory path containing installation images cannot contain commas or spaces. Click **Next** to continue.

Note: The installation images must be extracted before you set the image location.

11. The Pre-Installation Summary screen is displayed. Any installation errors are displayed here. Review the summary information on this screen and click **Back** to take corrective action, or click **Next** to continue.

Note:

- a. An error message is displayed on the Pre-Installation Summary screen if there is not enough space in the installation directory or the temp directory. Make sure that the directories you specified have enough free space.
 - b. If the Preinstallation Summary shows that a connection could not be established to the specified system, it might be because of a slow connection. Click **Back** to go to the previous page and click **Next** to try again. If the connection problem continues to exist, try installing at a time when there is less traffic on your network.
12. The Deployment screen displays the deployment status of each component and a progress indicator. When the deployment of each component has completed, click **Next** to display the Post Installation Summary page.

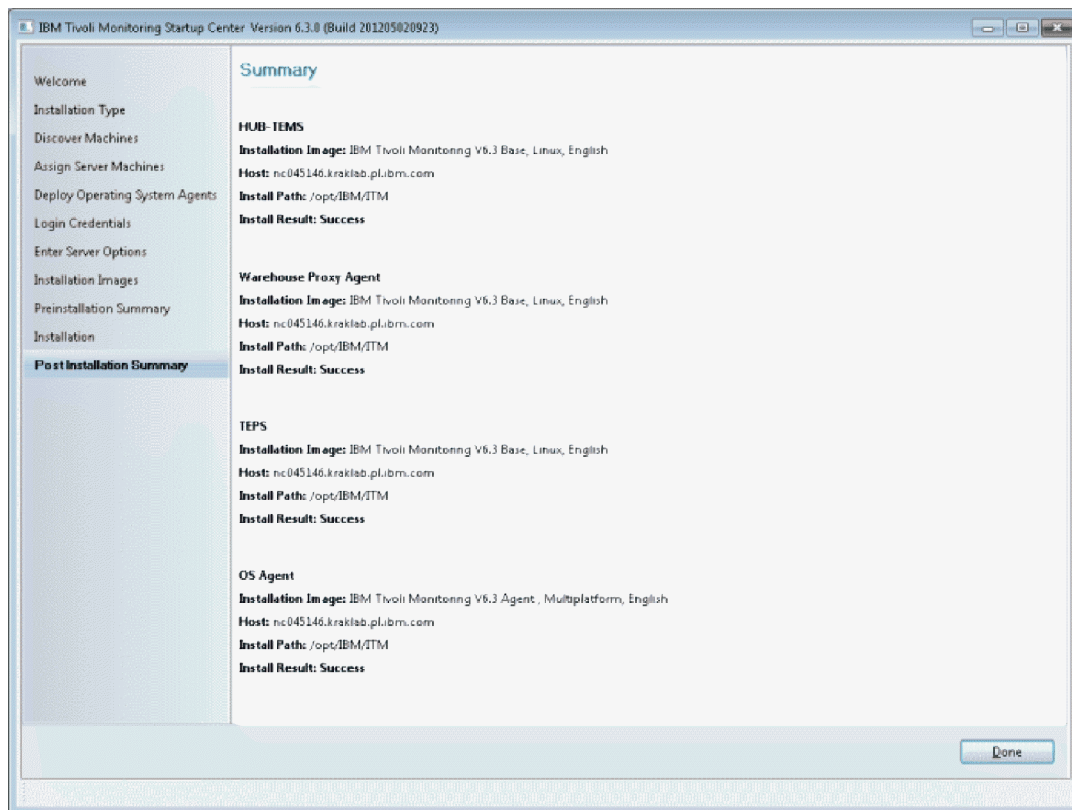


Figure 16. Tivoli Monitoring Startup Center Post Installation Summary window

Non-root user support

Be aware of the requirements when using a non-root user ID for Startup Center configuration.

- If you use a non-root user to log on to a target machine and get a connection failed error, it might be because the password provided has expired. Ensure that passwords have not expired before you attempt to install components by using the Startup Center.
- If you install a monitoring agent on a UNIX computer as a non-root user, the file permissions are initially set to a low level. Additional steps must be performed to complete the setup. The UpdateAutoRun.sh and SetPerm scripts must be executed as the root user on the target machines. For more information, see the procedure in “Postinstallation steps for nonroot installations” on page 270.
- For Windows machines, any non-root user must be a member of the administrator group.
- If you select DB2 as your Tivoli Enterprise Portal Server database, you must reconfigure the Tivoli Enterprise Portal Server independently by using root authority. Complete the following steps to reconfigure the Tivoli Enterprise Portal Server:
 1. Log in as a root user and check that your installation path has the correct permissions. If the installation path belongs to a non-root user, change the access permission level to 755. For example, if the candle home directory is /home/tester/ITM, the existing home and tester folders must have permission 755. Use the following command to change the permission level:


```
chmod 755/home/tester
```

where:

755

Is the new permission `rwxr-xr-x`.

/home/tester

Is the folder where IBM Tivoli Monitoring is installed.

2. Use the CLI tool or the Manage Tivoli Enterprise Monitoring Services window to reconfigure the Tivoli Enterprise Portal Server:
 - In the command-line interface, run the following command:
`<Candle_Home>/bin/itmcmd config -A cq`
 - In the Manage Tivoli Enterprise Monitoring Services, right-click **Tivoli Enterprise Portal Server** and click **Configure**.
3. The required DB2 user and database parameters have already been set up in the Startup Center. Complete the configuration process by using the default values provided.
4. Use the following commands to add a new group, such as **itmgroup**, and add both the non-root user and DB2 administrator user to the group:

```
groupadd itmgroup
usermod -a -G itmgroup tester
usermod -a -G itmgroup db2inst1
```

where:

tester

Is the non-root user.

db2inst1

Is the DB2 administrator user.

5. After you have made your changes, use the following command to run the secureMain utility before you start the Tivoli Enterprise Portal Server:

```
./secureMain -g itmgroup lock
```

Note: Whenever you reconfigure the Tivoli Enterprise Portal Server, you should run the secureMain utility before you restart the Tivoli Enterprise Portal Server.

For more information on the secureMain utility, see Appendix G, “Securing your IBM Tivoli Monitoring installation on Linux or UNIX,” on page 925.

6. Restart the Tivoli Enterprise Portal Server as a non-root user.
- For Tivoli Data Warehouse: Because the non-root user cannot create a DB2 user or reset a DB2 user's password, you must perform these tasks manually as a root user. You then add the user to the DB2 administrator group and re-configure the Tivoli Data Warehouse.

Default values taken by Startup Center components

This section contains the default values that are used for all components by the Tivoli Startup Center.

By default, the Tivoli Monitoring Startup Center names the hub Tivoli Enterprise Monitoring Server as `machinename_TEMS`. (If the Manage Tivoli Enterprise Monitoring Services is opened, or the `kincinfo -r` command is executed, the name 'TEMS1' is displayed. This is the name of a service under which the Tivoli Enterprise Monitoring Server is registered on Windows systems.) The remote Tivoli Enterprise Monitoring Server is named `machinename_RTEMS`. The following default values are used for all components:

INSTALL_ENCRYPTION_KEY=IBMTivoliMonitoringEncryptionKey

Encryption key.

SEED_TEMS_SUPPORTS=true

Seed this support on the monitoring server.

DEFAULT_DISTRIBUTION_LIST=NEW

Default seeding option for the *situation distribution definition* is NEW.

FIREWALL=NO

The Startup Center component does not connect to the Tivoli Enterprise Monitoring Server through a firewall.

PRIMARYIP=none

The primary IP address is not provided.

FTO=NO

Does not configure a connection for a secondary Tivoli Enterprise Monitoring Server.

NETWORKPROTOCOL=ip.pipe

The connection mode is IP_PIPE.

IPPIPEPORTNUMBER=1918

The default IP_PIPE port number is 1918.

The following default values are used for the Warehouse Proxy Agent:

- KHD_DB2_JDBCDRIVER=com.ibm.db2.jcc.DB2Driver
- KHD_BATCH_USE=true
- KHD_DB_COMPRESSION=false
- KHD_WAREHOUSE_TEMS_LIST=
- KHD_SERVER_Z_COMPRESSION_ENABLE=false
- KHD_SERVER_DIST_COMPRESSION_ENABLE=true

The following default value is used for the Tivoli Enterprise Portal Server:

- DB2ATTR=CONNECTION_LIMIT32

Remote Execution and Access

The Startup Center uses Remote Execution and Access (RXA) to establish a connection with a remote computer. RXA is an IBM developer toolkit that provides classes and methods to create an application that can establish a connection with a remote computer, log on to a remote computer, run commands and scripts on a remote computer (including installation and uninstallation), and manipulate the remote computer's files and directories.

RXA does not require the remote computer to have a software agent, such as the Tivoli management agent, installed on the remote machine. Instead of relying on a remote agent for communications with a remote computer, RXA provides access to the remote machine by using rsh, rexec, SSH, Windows (Server Message Block/Common Internet File System), and AS/400® host server protocols.

The default connection timeout setting for RXA in the Startup Center is 180000 milliseconds. You can alter this setting by changing the CONNECTION_TIME_OUT setting located here:

- **Windows** %USERPROFILE%\STARTUP\workspace\userdata\custom.properties
- **Linux** \${HOME}/.STARTUP/workspace/userdata/custom.properties

You must restart the Startup Center after you change this setting.

Windows Targets

Some RXA operations rely on VBScript and Windows Management Instrumentation (WMI) calls to execute scripts on Windows targets. If the Windows Scripting Host (WSH) or the WMI service is disabled on the target, or if VBScript is otherwise disabled, some WindowsProtocol methods will not work.

If you intend to access Windows targets by using the SMB protocol over NetBIOS, which is determined by setSMBTransportType(), then port 139 or the port specified by setNetBIOSPort(), must not be blocked by firewalls or IP security policies. The Enable NetBIOS over TCP/IP must also be selected in the Control

Panel settings for the machine's network connections properties (**Control Panel → Network and Dial-Up Connections → <some connection> → Properties → Internet Protocol (TCP/IP) → Advanced → WINS → Enable NetBIOS over TCP/IP**).

Consult the documentation for your firewall to determine that these ports are not blocked for inbound requests.

To determine if security policies are blocking these ports, click **Start → Settings → Control Panel → Administrative Tools**. Depending on whether your policies are stored locally or in Active Directory, the next steps are as follows:

- Locally stored policies: **Administrative Tools → Local Security Policy → IP Security Policies on Local Computer**.
- Policies stored in Active Directory: **Administrative Tools → Default Domain Security Settings → IP Security Policies on Active Directory**.

Examine the IP security policies and edit or remove filters that block the ports listed above. Table 24 lists the ports reserved for NetBIOS. Ensure that all ports currently used by RXA are not blocked.

Table 24. NetBIOS Reserved Ports

Port number	Use
135	NetBIOS Remote procedure call. At this time, RXA does not use this port.
137	NetBIOS name service.
138	NetBIOS datagram. At this time, RXA does not use this port.
139	NetBIOS session (file/print sharing).
445	CIFS (On XP and Win2K).

A utility program (testconn) is provided in the RXA .zip and JAR files. It can be used to determine whether a remote Windows target is configured to run Server Message Block protocol on top of NetBIOS (NetBIOS over TCP/IP) by using port 139, or whether the target is configured to run SMB on top of TCP/IP (without the NetBIOS layer) by using port 445. The target machine must have the Remote Registry service started (which is the default configuration) in order for RXA to connect to the target machine. A utility program (getregkey) is provided in the RXA zip and JAR files. It can be used to determine whether a Windows target (local or remote) has Remote Registry management enabled.

RXA requires access to the hidden remote administrative disk share for access to the system %TEMP% and other directories. Access to the Interprocess Communications share (IPC\$) is also required for RXA to access remote registries. Before you access the Interprocess Communications share (IPC\$), make sure the Server service is started (**Control Panel → Administrative Tools → Services → Server**). You can use the testconn.exe utility (in the \diagtools directory) to verify that the administrative share is accessible. RXA requires Simple File Sharing to be disabled. The next section details information specific to operating systems.

Disabling User Account Control to facilitate RXA

You might have to disable the User Account Control to enable RXA to connect to your Windows operating system.

Windows XP: Windows XP systems must have Simple File Sharing disabled for RXA to work. Simple Networking forces all logins to authenticate as "guest". A guest login does not have the authorizations required for RXA to function.

To disable Simple File Sharing, start Windows Explorer and click **Tools → Folder Options**. Select the **View** tab and scroll through the list of settings until you find **Use Simple File Sharing**. Remove the check mark next to **Use Simple File Sharing**, and then click **Apply** and **OK**.

Windows XP includes a built-in firewall called the Internet Connection Firewall (ICF). By default, ICF is disabled on Windows XP systems. Windows XP Service Pack 2 includes the Windows ICF set to ON by default. If either firewall is enabled on a Windows XP or Vista target, the firewall blocks access by RXA. On XP Service Pack 2, you can select the **File and Printer Sharing** box in the **Exceptions** tab of the Windows Firewall configuration to allow access.

Windows 2003: Windows 2003 systems must have Simple File Sharing disabled for RXA to work. Check that the firewall settings are the same as outlined for Windows XP above.

Windows Server 2008: On Windows Server 2008 you might need to disable User Account Control if your account is not a domain user account. See the section on Windows Vista to learn how to disable User Account Control.

Windows Vista: The new User Account Control feature in Windows Vista requires users to perform several steps before RXA applications can communicate with Vista targets. If you have a domain user account, ensure that the local and the target machine are both members of a Windows domain.

If you are a member of a local administrators group and you use a local user account, complete the three steps below to be able to perform administrative tasks on the target machine:

1. Enable the built-in Administrator account and use it to connect to the target system. To enable the built-in Administrator account, open the Windows Control Panel and click **Administrative Tools → Local Security Policy → Security Settings → Local Policies → Security Options**. Then double-click **Accounts: Administrator account status** and select **enable**.
2. Disable User Account Control if a different Administrator user account is to be used to connect to the Vista target. To disable User Account Control, open the Windows Control Panel and click **Administrative Tools → Local Security Policy → Security Settings → Local Policies → Security Options**. Then double-click **User Account Control: Run all administrators in Admin Approval Mode** and select **disable**. Changing this setting requires a system restart.
3. Disable User Account Control when you administer a workstation with a local user account (Security Account Manager user account). Otherwise, you cannot connect as a full administrator and cannot complete administrative tasks. To disable User Account Control, complete the following steps:
 - a. Click **Start**, click **Run**, type `regedit`, and then press **ENTER**.
 - b. Locate and then click the following registry subkey:
`HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System`
 - c. If the `LocalAccountTokenFilterPolicy` registry entry does not exist, follow these steps:
 - 1) On the Edit menu, point to New, and then click **DWORD Value**.
 - 2) Type `LocalAccountTokenFilterPolicy`, and then press **ENTER**.
 - 3) Right-click **LocalAccountTokenFilterPolicy**, and then click **Modify**.
 - 4) In the Value data box, type 1, and then click **OK**.
 - 5) Restart your computer.

Alternatively, you can modify the registry entry manually by typing the following command at a command prompt:

```
cmd /c reg add HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\
system /v LocalAccountTokenFilterPolicy /t REG_DWORD /d 1 /f
```

Windows 7: On Windows 7, the default startup type for the Remote Registry service is manual. The Remote Registry service must be running to enable RXA.

To check whether the Remote Registry service is enabled and started:

1. Go to **Start**.
2. In the Start Search box, type `services.msc`. Press **ENTER**.
3. When Microsoft Management Console starts, in the console pane, ensure that the service status is: started. If not, right-click **Remote Registry**, and click **Start**.

To avoid problems with the manual startup, it is advisable to set the Remote Registry service startup type to automatic. To automatically start the service after the server starts:

1. Right-click **Remote Registry** and select **Properties**.
2. In the Startup type option, choose **Automatic**.
3. Click **Apply** and **OK**.

When the system starts up, Remote Registry automatically starts.

Windows Vista FDCC (Federal Desktop Core Configuration): With Windows Vista FDCC custom security settings, it is not possible to connect to this operating system by using RXA.

Complete the following steps on Windows Vista FDCC to enable RXA to connect to the operating system:

1. Allow File and Printer Sharing with the Firewall by enabling the inbound file and printer exception by using the Local Group Policy Editor:
 - a. Go to **Start**.
 - b. In the **Start Search** box, type `gpedit.msc`. Press **ENTER**.
 - c. Go to: **Local Computer Policy** → **Computer Configuration** → **Administrative Templates** → **Network** → **Network Connections** → **Windows Firewall** → **Standard Profile** and enable Windows Firewall: **Allow inbound file and printer sharing exception**
2. Turn off the User Account Control.
3. Start the Remote Registry service.

Chapter 6. Post-deployment phase

There are some important steps that need to be taken to ensure your environment stays up and running and healthy.

Now that you have your Tivoli Monitoring environment in production, there are some important steps that need to be taken to ensure your environment stays up and running and healthy. This includes applying routine maintenance and performing periodic health checks. The post-deployment phase is divided into two sections: “Applying maintenance” and “Maintaining an efficient monitoring environment” on page 121.

Applying maintenance

This section outlines the planning and implementation steps necessary to install maintenance in your Tivoli Monitoring environment.

Routine maintenance is outlined in the following sections:

- “Planning an upgrade”
- “Upgrade steps”
- “Post-upgrade health check” on page 120

Planning an upgrade

Before upgrading your infrastructure components follow the checklist to plan your upgrade.

Use the following checklist to plan your upgrade:

- Check that the plan is in place for upgrading the environment (upgrade the environment incrementally). Follow a formal change management plan for Tivoli Monitoring upgrades and include, at minimum, both a deployment and tested backout plan.
- Download the correct upgrade media, *not* the fresh installation media.
- Back up all Tivoli Monitoring core components such as monitoring server and portal server.
- Carefully review the appropriate Fix Pack Readme and Documentation Addendum for any prerequisites.

Attention: Before upgrading your infrastructure components and beginning the upgrade process, perform a cold backup of your hub monitoring server, portal server, portal client, Warehouse Proxy Agents, Summarization and Pruning Agents, and remote monitoring server. Back up the following key components:

- Portal server database
- Warehouse database
- Full system backups and file system backups for installed Tivoli Monitoring components

Upgrade steps

When you are performing an upgrade, ensure that you install your IBM Tivoli Monitoring components in the correct order.

About this task

When performing upgrades read this *Installation and Setup Guide* or the supplied fix pack readme carefully. Perform your installation in the following order:

Note: This order might vary depending on the content of the release and the fix pack.

Procedure

1. Event Synchronization
2. Warehouse including Warehouse Proxy Agent and Summarization and Pruning Agent
3. Hub Tivoli Enterprise Monitoring Server
4. Remote Tivoli Enterprise Monitoring Server
5. Tivoli Enterprise Portal Server
6. Run any scripts necessary to update the Warehouse schema
7. Tivoli Enterprise Portal desktop client
8. Update monitoring agents

In some cases, when you upgrade the infrastructure components you will also be upgrading the monitoring agents on those servers. It is okay to upgrade those monitoring agents at that time.

Self-describing monitoring agents is a feature that integrates the installation of an agent with the dispersal and installation of associated product support files throughout your Tivoli Monitoring infrastructure.

“Self-describing agent installation” on page 394

Product support files for IBM Tivoli Monitoring agent applications must be installed at each IBM Tivoli Monitoring server component. The self-describing agent feature combines the installation of an agent with the automatic dispersal and installation of the associated product support files throughout your IBM Tivoli Monitoring infrastructure.

“Installing product maintenance” on page 370

The installation procedures in this chapter contain instructions for new installations. Follow the same procedures for upgrading or updating an existing installation. Fix packs (that is, product maintenance) use the same installer as a pristine installation except that in this case, the installer runs only on a system that is already installed and configured. This is true when installing both generally available fix packs and interim fixes.

Post-upgrade health check

Use the following checklist for your post-upgrade health check.

- Check if the predefined and user-created workspaces that were present prior to the upgrade are still in place.
Log in to the portal server by using the portal browser or desktop client and browse through the different workspaces for each product type.
- If you enabled the self-describing agent capability for your upgrade, run the **tacmd listappinstallrecs** to ensure that all products upgraded by the self-describing agent feature completed the installation without errors. Check the Audit Log messages for the Tivoli Enterprise Monitoring Server and Tivoli Enterprise Portal Server to ensure there are no self-describing agent errors related to the latest product upgrades.
- Check if the predefined and user-created situations that were present prior to the upgrade are still in place.
Log in to the portal server by using the portal browser or desktop client and browse through different situations in the Situation Editor or run the **tacmd listsit** command.
- Check if all the catalogs are the same on each monitoring server (hub and remote). Try one of the following two approaches:
 - Run grep on each monitoring server (hub and remote). For example:

```
grep @ * | awk '{print $2, $3, $4, $5}' | sort
```
 - Run the following SQL against each monitoring server in a portal server view:

```
"SELECT APPL_NAME, TIMESTAMP FROM SYSTEM.SYSAPPLS AT ('REMOTE_TEMS') ORDER BY APPL_NAME"
```
- Check if the depots populated on each monitoring server (hub and remote) are the same.

Run these commands from the hub monitoring server.

- `tacmd viewdepot`
- `tacmd viewdepot -j remote_tems`

- Check if the warehouse data is visible through the workspace views, meaning the portal server still has the correct connection to the warehouse database.

Select the attribute group for which history collection is enabled by checking that view and making sure the data can be pulled for more than 24 hours.

- Check if the agents are online and connected to the expected remote monitoring server.

Run the **tacmd listsystems** command.

- Check if the situations are firing and events are being forwarded to Tivoli Enterprise Console or OMNIbus.

Run the command on the Tivoli Enterprise Console server using **wtdump** or drag the Tivoli Enterprise Console icon to any view to view the events.

- Check if the historical configuration is active.

Log in to the portal server by using the portal browser or desktop client and click **History Configuration**. Browse through the desired attribute groups to see if they are still active.

Or you can run this query: "SELECT NODEL, OBJNAME, LSTDATE FROM 04SRV.TOBJACCL WHERE OBJNAME LIKE 'UADVISOR*'"

- Check if the Warehouse Proxy Agent and Summarization Pruning Agent are correctly started, meaning the agents made successful connections to the warehouse database.

You can examine the WAREHOUSELOG table to see the last updates by each attribute group. See sample below:

```
SELECT ORIGINNODE AS "Agent Hostname", OBJECT AS "Attribute Group",  
EXPORTTIME AS "Export Time", ROWSRECEIVED AS "Received Rows",  
ROWSINSERTED AS "Inserted Rows", ROWSSKIPPED AS "Skipped Rows",  
ERRORMSG AS "Error Message" FROM WAREHOUSELOG
```

Note to Windows users: If you attempt to run a tacmd CLI command and either the Embedded Java Runtime or the User Interface Extensions are not available on the node where you invoke the command, you will receive the error shown in Figure 52 on page 267. If this happen, complete the procedure outlined in “Installing the Embedded Java Runtime and the User Interface Extensions” on page 267 and retry the tacmd command.

Maintaining an efficient monitoring environment

This section covers the daily, weekly, monthly, and quarterly routine health checks on the Tivoli Monitoring enterprise environment.

- “Daily health checks” on page 122
- “Weekly health checks” on page 123
- “Monthly health checks” on page 123
- “Quarterly health checks” on page 123

By performing these routine procedures in addition to your daily, weekly, monthly and quarterly health checks, you ensure that your Tivoli Monitoring environment continues to run smoothly.

- Run the `taudit.js` tool which can be found in the IBM Integrated Service Management Library by searching for “Web SOAP scheduled reporting tools” or navigation code “1TW10TM0U”. This tool provides an overall status of the environment. Run this tool every day.
- Take a monitoring server backup every 24 hours in early stages and then move to weekly backups. If you have snapshot software, you can take backups with the monitoring server or portal server, or both online. Otherwise, shutdown the monitoring server and portal server before taking a backup. Test these backups after you first develop the process and at least twice a year thereafter by restoring to a

monitoring server in your test environment to ensure the backups are successfully backing up the data you need to restore your production monitoring server in the event of an outage or need for rolling back to a previous state.

- Make sure the portal server database backup is in the plan and is being made daily as the environment is being rolled out and then weekly as the environment matures and less frequent changes are made to the environment. Test these backups after you first develop the process and at least twice a year thereafter by restoring to a portal server in your test environment to ensure the backups are successfully backing up the data you need to restore your production portal server in the event of an outage or need to rollback to previous state.
- Make sure the DB2 warehouse backup is in the plan and is being made weekly. The reason you need to do this weekly is because of huge database size.
- Check daily that the warehouse agent is performing by looking at the warehouse logs (hostname_hd_timestamp-nn.log).
- Check daily that the Summarization and Pruning Agent is performing by looking at the (hostname_sy_timestamp-nn.log) logs.
- Check the monitoring server (hostname_ms_timestamp-nn.log) and portal server logs (hostname_cq_timestamp-nn.log) for any obvious errors and exceptions.
- Check that there are no monitoring servers overloaded with agents. One way to do this is by checking the "Self-Monitoring Topology" workspace, which has a "Managed Systems per TEMS" view showing the number of agents reporting to each monitoring server.
- For DB2, run the **REORGCHK** and **RUNSTATS** on the warehouse database daily.
- For DB2, run the **REORGCHK** and **RUNSTATS** on the portal server database weekly.
- Check that events are reaching the Tivoli Enterprise Console server or OMNIBus.
- Check that all the fired situations are answered with a response and are not in open state for a long period of time.
- Check that the Tivoli Enterprise Monitoring Server can communicate to each of the monitoring agents. The easiest way to test this is to run the `taudit.js` tool (mentioned above).
- Check the core components process memory and CPU usage and that you have situations created to monitor them.

Daily health checks

Use the following list to perform daily health checks.

Tasks

- Make sure all of the systems from the day before are still online. A quick look at the Managed System Status workspace shows the status of each managed system. If you encounter managed systems that are offline investigate them individually.

There are several reasons why a managed system can go offline. The agent might have gone offline for some reason, there may be communication problems between the agent and the monitoring server that it is connected to or the agent was decommissioned. In any case the cause of the problem must be found and addressed. Run a script every morning that provides a report on ONLINE and OFFLINE systems, `taudit.js` can be used for this purpose.

- You might find situations that are in open status that have not been addressed (acknowledged). Determine if the problem reported by the situation is valid. Determine if there is really a problem or is it a case where the situation does not have the correct thresholds and is producing a false positive that is being ignored by the monitoring staff. Make sure your situations are reflecting real events, which helps train the monitoring staff to react to each event that goes true in the Tivoli Monitoring environment.
- If you have decided to collect historical data and are using the Tivoli Data Warehouse, make sure the Warehouse Proxy Agent and Summarization and Pruning Agents are up and running. Check the logs for both to make sure the agents are collecting and summarizing data on the intervals you have set. To

find a solution that allows you to monitor the warehouse activity to ensure that it is functioning properly, search for “Data Warehouse DB activity” or navigation code “1TW10TM1X” in the IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>).

- Spot-check the workspaces for several different monitoring agent types to make sure report data is being returned.

Weekly health checks

Use the following list to perform weekly health checks.

Tasks

- Include all of the items in the daily health check.
- Check system backups. Have a mechanism in place for backing up your core Tivoli Monitoring components in the event of recovery. The portal server, monitoring server, and warehouse must be backed up on a regular interval. That interval must be decided by you, but make it no less than a weekly backup. Daily backups of the portal server and monitoring server databases are highly recommended because of their constant change.
- Check warehouse data. Make sure you are collecting the last week’s worth of data and it is correct. You can accomplish this by using the portal client to run the weekly summarization reports. The summarized data returned must accurately reflect the past week’s worth of summarized data.

Monthly health checks

Use the following list to perform monthly health checks.

Tasks

- Include all of the checks from the daily and weekly item checklist.
- If you are collecting historical data and storing it in the warehouse make sure your monthly summarization is correct. Validate this by running reports from the portal client to ensure you have the correct monthly summarized data.
- Check the list of managed systems deployed in your Tivoli Monitoring environment. Take note of their maintenance levels. Check with IBM Software Support or your IBM account representative to see if new fix packs and interim fixes are available. If so, determine what has been fixed so you can decide if you want to deploy the patches to your environment or just wait until the next major fix pack.
- Once again check your situation thresholds to make sure you don’t have false positive events. In large user environments there are many factors that can have an affect on how a system performs. The change in performance in any system can change the way Tivoli Monitoring reports status for any given system. Make sure the events active in Tivoli Monitoring are real.
- Take inventory of the systems being managed by Tivoli Monitoring. There might be a need to deploy additional agents on new systems or systems where new applications have been added.
- Assess the capacity of the infrastructure systems for resource CPU, memory and disk utilization to continually plan for overall workload balancing. As new versions of applications are introduced into the environment, their affect on resources typically change. This ongoing effort helps ensure the correct hardware is in place. Confirm the number of agents connected to each remote monitoring server to ensure that you have not exceeded the recommended limit of 1500 agents.

Quarterly health checks

Use the following list to perform quarterly health checks.

Tasks

- Include all of the checks from the daily, weekly, and monthly checklist.
- Discuss the usage of the historical data with the end users who rely on the data for various reasons. You might find they are not looking at some of the data being collected and no longer need it. In this

case, turn off the collection so you are not using unnecessary resources. The same holds true for the reverse. There might be data missing that is needed and historical collection can then be activated for the necessary information.

- Check with IBM Software Support or your IBM account representative for fix packs on all IT components. Regular maintenance fix packs for each component are typically delivered on a quarterly basis. Look through the Readme files and decide if you feel it is necessary to install the fix pack. Install the latest fixes for any of the components.

Chapter 7. Preparing for installation

The sections in this chapter provide an overview of the installation process and information to help you prepare to install your IBM Tivoli Monitoring environment.

The following table provides an overview of the steps required to fully install and deploy an IBM Tivoli Monitoring environment.

Table 25. Installation and configuration steps

Step	Where to find detailed information
Assess your monitoring needs to determine the best deployment of IBM Tivoli Monitoring components.	Chapter 2, "Pre-deployment phase," on page 19
Ensure you have the required hardware and software.	Chapter 3, "Hardware and software requirements," on page 75
Gather any information required for successful installation (such as DB2 user information and security specifications).	"Specific information to have ready" on page 126 Appendix A, "Installation worksheets," on page 833
Install the Tivoli Enterprise Monitoring Server.	"Installing and configuring the hub Tivoli Enterprise Monitoring Server" on page 196
Install the Tivoli Enterprise Portal Server.	"Installing the Tivoli Enterprise Portal Server" on page 226
Install the monitoring agent software.	"Installing monitoring agents" on page 261 "Self-describing agent installation" on page 394
Setup event forwarding to Netcool/OMNIBus or Tivoli Enterprise Console.	Chapter 30, "Setting up event forwarding to Netcool/OMNIBus," on page 697 or Chapter 31, "Setting up event forwarding to Tivoli Enterprise Console," on page 799
Install the portal desktop client on any system where you want to use it.	"Installing the Tivoli Enterprise Portal desktop client" on page 271
Start the portal client to verify that you can view the monitoring data.	"Starting the Tivoli Enterprise Portal client" on page 366
Install and setup the Tivoli Data Warehouse	Chapter 19, "Tivoli Data Warehouse solutions," on page 509
Install and setup a dashboard environment (optional)	"Required software and memory requirements for a dashboard environment" on page 91 "Using IBM Installation Manager to install or update IBM Tivoli Monitoring components for a dashboard environment" on page 299 "Setting up your dashboard environment" <i>IBM Tivoli Monitoring Administrator's Guide</i>
Install and setup a reporting environment with Tivoli Common Reporting (optional)	"Required software and memory requirements for a reporting environment that uses Tivoli Common Reporting" on page 93 and the <i>Tivoli Common Reporting</i> chapter in the <i>IBM Tivoli Monitoring Administrator's Guide</i> .

Table 25. Installation and configuration steps (continued)

Step	Where to find detailed information
Install and setup OSLC linked data integration using the OSLC Performance Monitoring service provider (optional)	<p>“Required software and memory requirements for using the OSLC Performance Monitoring service provider” on page 94</p> <p>“Installing and configuring Tivoli Enterprise Monitoring Automation Server” on page 219</p>

If you are upgrading from IBM Tivoli Monitoring V6.1, see Chapter 8, “Upgrading from a previous installation,” on page 145 before installing any IBM Tivoli Monitoring components.

If you are upgrading from Tivoli Distributed Monitoring or IBM Tivoli Monitoring V5.x to V6.3, see *IBM Tivoli Monitoring Upgrading from Tivoli Distributed Monitoring*.

If you plan to use firewalls in your environment, see Appendix C, “Firewalls,” on page 861 for an overview of the IBM Tivoli Monitoring implementation of firewalls.

Specific information to have ready

This section contains the information that you must have ready for your Tivoli Monitoring installation.

During installation, you must supply the following information:

- Name of the monitoring server you are installing or to which the agent will connect
- Host name for the computer where you are installing the product (a monitoring server or one instance of an agent)
- Whether the monitoring server being installed or being connected to is configured as a hub or remote monitoring server
- Hub monitoring server host name
- Port number

Use the worksheets in Appendix A, “Installation worksheets,” on page 833 to collect this information for each component that you want to install.

Information to gather for event forwarding

You need the following additional information to successfully install and configure event forwarding and synchronization between a hub Tivoli Enterprise Monitoring Server and either Netcool/OMNIBus or IBM Tivoli Enterprise Console.

- Host names, user IDs, and passwords for the monitoring servers that you want to receive events from.
- The amount of free space in your temporary directory on the server where Netcool/OMNIBus ObjectServer or Tivoli Enterprise Console are installed. The event synchronization component installation requires 200 MB of temporary space.
- Simple Object Access Protocol (SOAP or Web Services) information to send events to a monitoring server (the URL, the rate to send requests to the server).

By default, all Hub monitoring servers are configured as SOAP servers. If you did not change the SOAP configuration to make it unique for your environment, you can accept the default values during the installation.

If you did change this configuration, use the SOAP information unique to your configuration.

- For Netcool/OMNIBus, the host of the EIF probe and the port on which it is listening. For TEC, the host of the event server or servers to which events are being forwarded and the port on which it is listening.

- For TEC, event rule base information (either the name of a new rule base to create or the name of an existing rule base to use)

Note:

1. For a Windows event server, any existing rule base that you use must indicate a relative drive letter (such as C:\) as part of its associated path. To verify that your existing rule base contains a relative drive letter, run the following command from a bash environment on your event server:

```
wrb -lsrb -path
```

If the returned path includes something like *hostname:\Rulebase_directory*, with no drive letter (such as C:\), copy the ESsync3000Win32.exe file from the \TEC subdirectory of the IBM Tivoli Monitoring installation image to the drive where the rule base exists and run the installation from that file.

2. If you are using a Windows event server, if you have any rule base with an associated path that does not contain a relative drive letter and that has the Sentry2_0_Base class imported, copy the ESsync3000Win32.exe file from the \TEC subdirectory of the IBM Tivoli Monitoring installation image to the drive where the rule base exists and run the installation from that file.

To verify if you have any rule bases that have an associated path containing no relative drive letter, run the **wrb -lsrb -path** command as described in the previous note.

To determine if your rule bases have the Sentry2_0_Base class imported, run the following command against all of your rule bases:

```
wrb -lsrbclass rule_base
```

where *rule_base* is the name of the rule base.

Naming your monitoring server

You must decide how the monitoring servers are to be named.

In general, create names that are short but meaningful within your environment. Use the following guidelines:

- Each name must be unique. One name cannot match another monitoring server name for its entire length. (For example, "ibm" and "ibmremote" are unique and permitted.)
- Each name must begin with an alpha character. No blanks or special characters ("\$#@") can be used.
- Each name must be between 2 and 32 characters in length.
- Monitoring server naming is case-sensitive on all platforms.

Choose between IPv6 and IPv4

You can now enable IP version 6 (IPv6) communication between any two IBM Tivoli Monitoring components; possible configurations include communications between the portal server and the hub monitoring server, a remote monitoring server and the hub, or an agent and a hub.

IPv6 is not supported for sending EIF events from the hub monitoring server or from monitoring agents.

To use this capability, your IBM Tivoli Monitoring environment must be configured and enabled for IPv6 communication. IPv6 communication over IPv4-only networks is not supported.

Before components can be enabled for IPv6 communication, the following requirements must be met:

1. The host on which the components are located must be enabled for IPv6.
2. If the component needs to communicate using IPv6 with some components and IPv4 with others, the host must be enabled for dual-stack operation.

Note: A dual-stack host provides two discrete network layers. The term *stack* here refers to the protocol stack, a suite of protocols used in computer networking software.

3. The host must have DNS or hosts file entries for both IPv4 and IPv6 addresses. The host must also be configured to resolve host names and IP addresses from the DNS or from the hosts file.
4. The network path between the two components must be enabled for IPv6 traffic. Tunneling of IPv6 traffic over IPv4 networks is not supported.

Table 26 shows the supported combinations of IPv6 with IPv4 across the various IBM Tivoli Monitoring components.

Table 26. Supported IBM Tivoli Monitoring configurations using the IPv6 communications protocol

Valid configurations	Portal client	Portal server	Hub monitoring server and automation server	Agents	Remote monitoring server	Agents (connected to the remote monitoring server)
IPv6 only	IPv6	IPv6	IPv6	IPv6	IPv6	IPv6
IPv6 with IPv4	IPv4	IPv4 or IPv6	IPv4 or IPv6	IPv4 ¹	IPv4 or IPv6	IPv4 ¹

Note:

1. All agents running on a computer must be configured to use the same protocol, either IPv4 or IPv6.
2. In scenarios where some agents are on IPv4 only computers or the network between the agents and the monitoring servers they report to is IPv4 only, these agents need to communicate with the monitoring servers over IPv4. The monitoring servers therefore may communicate with some agents over IPv4 and with others over IPv6.
3. The portal server does not support IPv6 on the Windows platform. If the portal server is on Windows, the browser and desktop clients need to communicate with it using IPv4.
4. Components do not operate in dual-stack mode on the Solaris platform. Components can be configured to communicate using either IPv4 or IPv6. Thus, if a hub server on a Solaris host is configured to use IPv6, the portal server, all remote servers, and all agents connecting to the hub must be configured to use IPv6 for communicating with the hub.
5. On HP-UX, patch PHNE_29445 is required for IPv6 support.
6. Components do not operate in dual-stack mode on the HP-UX HP9000 platform. Dual-stack mode is supported on the HP-UX Integrity platform.
7. On Linux computers, a minimum kernel level of 2.6 is required for IPv6 support.

Monitoring components, when installed and configured using the appropriate platform-specific configuration tools, are initially configured only for IPv4 communication on all platforms except z/OS (where your ICAT settings govern the protocol used). On all other platforms, you must perform supplemental configuration steps to reconfigure Tivoli Monitoring components to communicate using IPv6.

For more information, see Chapter 15, “Configuring IBM Tivoli Monitoring components for IPv6 communication,” on page 441.

Required order of installation or upgrade of IBM Tivoli Monitoring component products

This section contains information about the required order of installation or upgrade of IBM Tivoli Monitoring component products.

If any of the following products will be installed on the same computer as monitoring agents, they must be installed before the agent is installed:

- Hub Tivoli Enterprise Monitoring Server

- Remote monitoring server (if necessary)
- Tivoli Enterprise Management Agent Framework
- Tivoli Enterprise Portal Server
- Tivoli Enterprise Portal desktop client
- Dashboard Application Services Hub (a Jazz for Service Management component) ¹
- IBM Infrastructure Management Dashboards for Servers ¹
- Tivoli Authorization Policy Server ¹
- tivcmd Command Line Interface for Authorization Policy ¹
- Tivoli Enterprise Monitoring Automation Server ²

Note:

1. Install these components if you are planning to setup a dashboard environment.
2. Install this component if you are planning to use the Performance Monitoring service provider to integrate with the Jazz for Service Management Registry Services component and other products that support integration using OSLC.

In addition, these products must be installed on at least one computer before the agent can be properly configured. The appropriate Tivoli Enterprise Management Agent Framework is installed when an agent is installed.

Windows installation considerations

The following sections provide information about issues unique to Windows installations.

User authority

To install IBM Tivoli Monitoring on a Windows computer, you must have Administrator privileges on that computer. You must also run the IBM Tivoli Monitoring components as a user with Administrator privileges.

Windows 32 bit On Windows 64 bit (WOW64)

WOW64 is a subsystem of the Windows operating system that enables 32 bit Windows-based applications to run on 64 bit Windows.

WOW64 is a prerequisite for IBM Tivoli Monitoring. Beginning with Windows Server 2008 R2, WOW64 is an optional feature that you can uninstall. You must ensure that WOW64 is installed on your Windows computer before you install IBM Tivoli Monitoring. For more information, see <http://msdn.microsoft.com/en-us/library/windows/desktop/dd371790%28v=vs.85%29.aspx>.

32 bit versus 64 bit

If your site runs either Windows 2003 or 2008 on 64-bit x86-64 CPUs, you must decide whether to install the 32-bit operating system agent or the 64-bit agent.

The new 64-bit agent supports 64-bit operations and can coexist with other 32-bit agents in your Tivoli Monitoring environment.

Note:

1. Direct upgrade of the 32-bit Windows agent to the 64-bit agent is not supported. When upgrading a 32-bit agent from a prior release to the current release, you can upgrade only to the current 32-bit agent.
2. This new support does not extend to native 64-bit applications (including the operating system) running on the Itanium IA64 architecture.

3. This 64-bit support has not been extended to the IBM Tivoli Monitoring servers (the Tivoli Enterprise Monitoring Server and the Tivoli Enterprise Portal Server).
4. Support was added for a 64-bit Java Runtime Environment (JRE) to the V6.2.3 release of IBM Tivoli Monitoring. Prior releases supported 32-bit JRE only. The 64-bit JRE can only work on 64-bit machines. The ITM Installer does not allow a 64-bit JRE to be installed on a 32-bit machine. The installation summary dialog informs you of the JRE type being installed.

Installation using a Citrix client

If you are using a Citrix client to access the IBM Tivoli Monitoring installation program for Windows through Microsoft Windows Terminal Services, you must manually change Terminal Services to install mode before running the installation.

To change Terminal Services to install mode, run the **change user /install** command before starting the installation. After installation, run the **change user /execute** command to return Terminal Services to normal mode.

Memory leak issue on Windows Server 2008

You can monitor the private bytes of the Monitoring Agent for Windows OS agent process (kntcma.exe), either by using IBM Tivoli Monitoring or system tools like the Windows Task Manager. You might notice an increase in private bytes over time. After the Windows OS agent runs for several days, the computer might become unresponsive.

Microsoft hotfix KB949700 must be installed to resolve a memory leak issue on Windows Server 2008 systems. To download the hotfix, go to <http://support.microsoft.com/kb/949700>. Restart your computer after you have installed the hotfix.

Linux or UNIX installation considerations

The following sections provide information about issues unique to Linux and UNIX installations.

- “Changes in the behavior of the autostart scripts”
- “Create an IBM Tivoli account for installing and maintaining the installation directory” on page 134
- “Host name for TCP/IP network services” on page 134
- “Use of fully qualified path names” on page 135
- “Multiple network interface cards” on page 135
- “Installing into an NFS environment” on page 135
- “Installing into Solaris zones” on page 135
- “File descriptor (maxfiles) limit on UNIX and Linux systems” on page 137

Changes in the behavior of the autostart scripts

The behavior of the autostart scripts generated during installation and configuration on UNIX platforms has evolved.

- In V6.1 fix pack 3, the installation process produced an autostart script with only one entry using a generic CandleAgent start all command. Users modified this file as needed.
- In V6.1 fix pack 4, the installation process generated individual entries for each application in a particular installation, but the values captured in the file could not be overridden.
- In V6.1 fix pack 5, the multiple entries remain, and an override capability has been added.
- In V6.2.2 fix pack 2 and later, the multiple entries remain, and the override capability has been significantly enhanced.

The autostart script generated by an installation, upgrade, or configuration and named ITMAgentsN or rc.itmN (depending on the UNIX platform) contains an entry for each application in a particular installation. The entries look similar to this:

```
su - USER -c "ITM_Install_Home/bin/itmcmd agent start product_code"
```

Or:

```
su - USER -c "ITM_Install_Home/bin/itmcmd agent -o Instance start product_code"
```

Where:

USER

Is the ID that the application will be started as. By default, *USER* is the owner of the bin directory for the application. For the UNIX Log Alert agent, *USER* is the owner of the *ITM_Install_Home/PLAT/ul/bin* directory.

N Is an integer specific to each installation on a system.

ITM_Install_Home

Is the full path to the IBM Tivoli Monitoring version 6.x installation directory.

product_code

Is the two-character code for this application. See Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883 for a list of the codes for the common components and the base agents. See the product documentation for other product codes.

instance

Is the instance name required to start this application.

PLAT

Is the platform directory where the application is installed.

Components are started in the order listed in the autostart script. This order is based on the dependencies between components, rather than any logical sequence.

The *kcirunas.cfg* file was added to allow overrides to the default processing. The *kcirunas.cfg* file is delivered in the root directory of the installation media, in the same location as *install.sh*. During installation, this file is copied to the *ITM_Install_Home/config* directory (but is not overwritten if this file already exists). This file is provided as a sample file with each section commented out. You do not have to modify this file if you want the autostart script to be generated with the default processing.

For local installation usage, you may modify the *kcirunas.cfg* file in the root directory of the media if you want to use the same set of values for multiple installations on similar systems from this image. You may also modify the *kcirunas.cfg* file in the *ITM_Install_Home/config* directory if you want to use a specific set of values for each individual installation from this image.

For remote deployment usage, you can modify the *kcirunas.cfg* file in the root directory of the media. You can also modify the *kcirunas.cfg* file in the Tivoli Enterprise Monitoring Server depot after populating the depot from this image. If you have set the **DEPOTHOME** variable in the *tables/TEMS_NAME/KBBENV* file, you must use that value as the base when searching for the depot location. To determine if you have set **DEPOTHOME**, run the following commands:

```
cd ITM_Install_Home
DEPOTHOME=`find tables -name KBBENV -exec grep DEPOTHOME {} \; 2> /dev/null | cut -d= -f2`
echo $DEPOTHOME
```

If **DEPOTHOME** is not empty, run the following commands to locate *kcirunas.cfg* in the monitoring server depot:

```
cd ITM_Install_Home
DEPOTHOME=`find tables -name KBBENV -exec grep DEPOTHOME {} \; 2> /dev/null | cut -d= -f2`
find $DEPOTHOME -name kcirunas.cfg -print
```

If DEPOTHOME is empty, run the following commands instead:

```
cd ITM_Install_Home
find tables -name kcirunas.cfg -print
```

The file kcirunas.cfg contains a superset of the XML syntax and structure in the *ITM_Install_Home/config/HOST_kdyrunas.cfg* file (where *HOST* is the short hostname for this system) produced by remote configurations, such as remote deployment or Tivoli Enterprise Portal-based agent configuration.

The entries in kcirunas.cfg do not affect the actions performed for remote deployment, remote configuration, remote starting or stopping, or any Tivoli Enterprise Portal-initiated agent action. The entries in *HOST_kdyrunas.cfg* affect the generation of the reboot script. The entries in kcirunas.cfg also affect the generation of the reboot script, and they override any entries for the same component in *HOST_kdyrunas.cfg*.

The following is the default kcirunas.cfg file with all **<productCode>** entries commented:

```
<agent>

  <!--productCode-->ux</productCode-->
  <instance>
    <user>itmuser</user>
  </instance>

  <!--productCode-->ul</productCode-->
  <instance>
    <user>root</user>
  </instance>

  <!--productCode-->lz</productCode-->
  <instance>
    <user>itmuser</user>
  </instance>

  <!--productCode-->ud</productCode-->
  <instance>
    <name>db2inst1</name>
    <user>db2inst1</user>
  </instance>
  <instance>
    <name>db2inst2</name>
    <user>root</user>
  </instance>

  <!--productCode-->ms</productCode-->
  <instance>
    <name>HUB17</name>
    <user>itmuser</user>
  </instance>

  <!--productCode-->cq</productCode-->
  <instance>
    <user>itmuser</user>
  </instance>

  <!--productCode-->cj</productCode-->
  <instance>
    <user>itmuser</user>
  </instance>

</agent>
```

By default, each **<productCode>** section in the kcirunas.cfg file is disabled by making the product code a comment, such as **<!--productCode-->**. To activate a section, do the following:

1. Remove the comment indicator (the exclamation point, !) so that the `<!productCode>` item looks like `<productCode>`.
2. Copy a `<productCode>` section.
3. Rather than creating new sections from scratch, customize each `<productCode>` section, and activate it.

Commented, or *deactivated*, sections are ignored. Uncommented, or *activated*, sections for applications that are not installed are ignored. For agents that do not require an instance value, specify only:

```
<productCode>product_code</productCode>
<instance>
  <user>USER</user>
</instance>
```

For agents that do require an instance value, like the DB2 monitoring agent (product code ud), specify the *product_code*, *instance*, *user*, and *name*:

```
<productCode>ud</productCode>
<instance>
  <name>db2inst1</name>
  <user>db2inst1</user>
</instance>
<instance>
  <name>db2inst2</name>
  <user>root</user>
</instance>
```

Two items that are supported in the `kcirunas.cfg` file that are not supported in the `HOST_kdyrunas.cfg` file are the `<default>` section and the `<autoStart>` flag. The `<autoStart>` flag can be used in the `<default>` section and in the `<instance>` section. The `<default>` section is specified as follows:

```
<productCode>product_code</productCode>
<default>
  <user>db2inst1</user>
</default>

<productCode>product_code</productCode>
<default>
  <autoStart>no</autoStart>
</default>

<productCode>product_code</productCode>
<default>
  <user>db2inst1</user>
  <autoStart>no</autoStart>
</default>
```

The `<autoStart>` flag is specified as follows:

```
<productCode>product_code</productCode>
<default>
  <autoStart>no</autoStart>
</default>

<productCode>product_code</productCode>
<instance>
  <autoStart>no</autoStart>
</instance>
```

A section similar to the following can be used to not automatically start the default MQ Monitoring instance and to automatically start all other instances as the `mqm` user:

```
<productCode>mq</productCode>
<default>
  <user>mqm</user>
</default>
```

```
<instance>
  <name>None</name>
  <autoStart>no</autoStart>
</instance>
```

A set of sections similar to the following can be used to avoid automatically starting the set of installed agents and servers. You need one section for each agent or server component installed:

```
<productCode>product_code</productCode>
<default>
  <autoStart>no</autoStart>
</default>
```

Where *product_code* is the two-character product code for the individual agent or server component (See Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883).

Note:

1. Any changes made directly to the autostart script (ITMAgents*N* or rc.itm*N*, depending on the platform) will not be preserved and will be overwritten the next time you install, configure, or upgrade an application.
2. Any changes made to the AutoRun.sh script will not be preserved and will be overwritten the next time you apply higher maintenance.

Create an IBM Tivoli account for installing and maintaining the installation directory

Create an IBM Tivoli account for installing and maintaining the installation directory.

For best performance, follow these guidelines:

Note: This option does not apply to configuring the portal server on Linux systems where you may use a non-root IBM Tivoli Monitoring user ID to install the portal server. If you do, you must then use the root user ID to configure the portal server because the DB2 installation or administrator ID may lack the necessary privileges.

If, however, you use either the root user ID, the DB2 installation ID, or the DB2 administrator ID to install the portal server, you must use that same user ID to configure it. You can then use your non-root Tivoli Monitoring user ID to run the portal server.

- You can use any valid name.

You can install the IBM Tivoli Monitoring software as the root user, but you do not have to. If you do not install as a root user, you must follow the steps outlined in “Postinstallation steps for nonroot installations” on page 270 after you install any monitoring agents.

- Use the same user to install all components.
- If you are using NFS or a local file system, establish your installation directory according to the guidelines used in your environment.
- Only the Korn shell is supported for the execution of the installation and runtime scripts. Consider using the Korn shell as the default environment for your IBM Tivoli login account.

Host name for TCP/IP network services

TCP/IP network services such as NIS, DNS, and the /etc/hosts file must be configured to return the fully qualified host name (for example: HostName.ibm.com). Define the fully qualified host name after the dotted decimal host address value and before the short host name in the /etc/hosts.

On a machine that will be running the Tivoli Enterprise Portal Server, you should verify that the /etc/hosts file has only one line containing the short name of the host machine. Having multiple entries (that is different network cards) resolving to the same short name of the machine can result in the Tivoli

Enterprise Portal Server generating an error when either starting, reconfiguring, or running the `buildPresentation`, `InstallPresentation.sh`, or `migrate-export` scripts.

Use of fully qualified path names

Because of the wide variety of UNIX operating systems and possible user environments, use fully qualified path names when entering a directory during the installation process (do not use pattern-matching characters).

IBM scripts use the Korn shell. When a new process or shell is invoked, use of symbolic links, environmental variables, or aliases can potentially cause unexpected results.

Multiple network interface cards

When more than one network interface card (NIC) exists in the computer on which the monitoring server is installed, you need to identify which NIC to use when specifying the monitoring server name and host name.

Additionally, the host name of the system might not match the interface name, even when only one NIC exists. In either of these cases, to establish connectivity between the monitoring server and agents, you must specify an additional variable when configuring the monitoring server or agents. This variable is listed under the **Optional Primary Network Name** option in the configuration windows or during the installation.

If the host of the Tivoli Enterprise Portal Server has more than one NIC, you need to configure an additional interface for each one.

Installing into an NFS environment

IBM supports installing IBM Tivoli Monitoring in NFS environments. Using NFS, you can concentrate your software and data in a specific location, minimizing maintenance, administrative overhead, and disk space.

Although using NFS to support multiple hosts simplifies the maintenance of installed IBM Tivoli products, its use can impact performance. If you are installing into an NFS environment, consider the administrative savings to the possible impact on the performance of your network.

Consider the number of hosts that share a single installation directory, as well as the effects of network congestion and file system performance on the overall response time of your IBM Tivoli products.

NFS also has some trade-offs in how you manage your environment. While you can have your entire IBM Tivoli Monitoring in one place, there might be additional configuration required to define the use of specific products or processes in your installation directory. Will every product on every host system execute using the same configuration; or will you tailor the configuration to the particular environment?

Note: If installing from images on an NFS mount, the NFS mount needs **world** execute permissions to be accessible by the installation processes.

Installing into Solaris zones

There are limitations that must be taken into consideration when installing into zones that share resources.

Big local zones share no resources with other zones, so there are no limitations on installing into a Big Local zone. For local zones that share resources, there are the following limitations:

- Anything installed outside of `$CandleHome` must:
 - Be installed in the global zone, OR

- Grant local zone write access to the global zone directory tree
- Any devices used by an agent to obtain statistics must be linked to the Local zone.
- During installation of the Tivoli Enterprise Monitoring Server, the GSKit library can be installed from an alternate directory within the local zone. To accomplish this, edit the ms.ini file for the monitoring server, and add the extra library path to the LD_LIBRARY_PATH concatenation.

GSKit, the Global Security Toolkit, provides SSL (Secure Sockets Layer) processing within protocols such as SPIPE and HTTPS.

The following sections describe the types of zones and the installation possibilities for each.

Global Zone

The main administration zone. Both local and remote installation is possible. GSKit installs into the standard location for Solaris, and the links are created in the standard location for Solaris.

Big Local Zone

A local zone with no shared file system resources. Local and remote installation is possible. GSKit installs into the standard location for Solaris, and the links are created in the standard location for Solaris. These locations are local to this zone and not shared with any other zone.

Small Local Zone

A local zone with some shared file system resources. Local and remote installation is possible if /opt is not a shared resource. GSKit installs into the standard location for Solaris, but the links are created in the new default location of \$CANDLEHOME/gsklinks. These locations are local to this zone and not shared with any other zone.

Default Small Local Zone

A local zone with the default set of shared file system resources. This is a Small Local Zone with /sbin /lib /usr and /export directories shared with global zone. These directories are read-only file systems in local zone. Local and remote installation is possible. GSKit installs into the standard location for Solaris, but the links are created in the new default location of \$CANDLEHOME/gsklinks. These locations are local to this zone and not shared with any other zone.

Small Local Zone with /opt directory shared

Local and remote installation is not possible. Tivoli Monitoring installation always requires read-write access to the /opt directory. This is not only a GSKit issue. Even if CANDLEHOME is specified as the nondefault directory, read-write access to /opt/IBM/ITM/tmaitm6/links is still needed.

Note: In all supported Small Local Zone configurations, the Tivoli monitoring interactive command-line installation prompts you for the parent directory where the links will be created. For example, if you enter /tmp for the directory, the links will be created in the /tmp/usr/lib and /tmp/usr/bin directories. The default directory for this prompt is \$CANDLEHOME/gsklinks. During remote installation, the default directory is always used.

It is very difficult to predict all the possible shared-resource policies for small local zones and the possible side effects. It is the responsibility of the system administrator to create these policies without causing unintentional side effects between different zones and software installed.

Architecture and product codes

On UNIX and Linux operating systems, some IBM Tivoli Monitoring files are placed in directories whose paths include a code for the platform architecture.

You can find the codes for the supported architectures in the registry directory. **archdsc.tbl** contains the architecture codes and descriptions. You can find the product codes and descriptions for supported components and products in the same directory, in **proddsc.tbl**.

File descriptor (maxfiles) limit on UNIX and Linux systems

The monitoring server and Warehouse Proxy Agent can use a large number of file descriptors, especially in a large environment. On UNIX and Linux systems, the maximum number of file descriptors available to a process is controlled by user limit parameters.

To display the user limits, run the following command:

```
ulimit -a
```

The "nofiles" parameter is the number of file descriptors available to a process. For the monitoring server process (kdsmain), the "nofiles" parameter should be set larger than the maximum number of agents that will be connecting to the monitoring server. If the monitoring server is unable to get file descriptors when needed, unexpected behavior can occur, including program failures. Consider increasing the value to 8000 file descriptors or more.

There are other user limit parameters that control how much data, stack and memory are available to a process. For large environments, consider increasing these memory-related user limit parameters for the monitoring server (kdsmain) process. Configuring the user limit parameters usually requires root access, and involves changing system startup files which are operating system specific. Consult the operating system manuals for information on how to configure the user limit parameters.

Installing into an existing installation

Installing components or agents in an existing CANDLEHOME or installation directory is supported as long as the user ID used to run the installation is always the same.

Installing components or agents in an existing CANDLEHOME or installation directory using different user IDs is not supported.

Security options

User IDs and passwords sent between Tivoli Management Services components are encrypted by default. Other communication between components can be secured by configuring the components to use secure protocols.

See "Communication between components" on page 138 for more information.

Access to the Tivoli Enterprise Portal (*authorization*) and tacmd commands that send requests to the portal server are controlled by user accounts (IDs) defined to the Tivoli Enterprise Portal Server. The hub Tivoli Enterprise Monitoring Server can be configured to validate, or *authenticate*, user IDs through either the local operating system registry or an external LDAP-enabled registry. Alternatively, authentication by an external LDAP registry can be configured through the Tivoli Enterprise Portal Server. If authentication is not configured through either the monitoring server or the portal server, no password is required to log on to the Tivoli Enterprise Portal. See "Authorization and authentication" on page 139.

Users that execute tacmd commands that send SOAP requests to the hub monitoring server, or user IDs that require direct access to the SOAP server, must be authenticated through the hub monitoring server. If user authentication is *not* enabled on the hub monitoring server, anyone can make requests to the SOAP Server. If user authentication *is* enabled on the hub, the SOAP Server honors only requests from user IDs and passwords authenticated by the local or external registry. If the type of access is specified for specific users, only requests from those users for which access is specified are honored. See "SOAP server security" on page 142. The user ID and passwords that are used to authenticate with the SOAP server must be 15 characters or less.

If you are using the Dashboard Application Services Hub with a monitoring dashboard application such as IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM or in custom dashboards, you should configure Dashboard Application Services Hub and Tivoli Enterprise Portal Server to use a central LDAP registry and enable single sign-on. This ensures that the portal server can authenticate dashboard users when they request monitoring data and Dashboard Application Services Hub forwards the request to the dashboard data provider component of the portal server. See “Single sign-on capability” on page 141. If you have previously enabled authentication through the hub monitoring server and want to use LDAP authentication and single signon with the portal server, see the *Enabling user authentication* chapter in the *IBM Tivoli Monitoring Administrator's Guide*.

Single sign-on should also be configured for the Tivoli Enterprise Portal Server if users will launch out of the Tivoli Enterprise Portal to other Tivoli Web-based or Web-enabled applications, or if they will launch into the Tivoli Enterprise Portal from other Web-based applications. Using single sign-on with a central LDAP registry allows users to move seamlessly between applications without having to re-enter their user IDs and passwords.

If you want the Performance Monitoring service provider to authenticate HTTP requests from OSLC clients, you must configure the service provider to use the Security Services component of Jazz for Service Management. Security Services is an optional Jazz for Service Management component that enables non-WebSphere based applications such as the Performance Monitoring service provider to participate in LTPA based single sign-on. See “Single sign-on capability” on page 141 for more details on using Security Services with the Performance Monitoring service provider.

Note:

1. The Tivoli Directory Server (TDS) LDAP client used by the Tivoli Enterprise Monitoring Server does not support LDAP referrals, such as those supported by Microsoft Active Directory.
2. The IBM Tivoli Monitoring Service Console enables you to read logs and turn on traces for remote product diagnostics and configuration. The Service Console performs user authentication using the native operating system security facility. This means that if you use the Service Console on z/OS, your user ID and password are checked by the z/OS security facility (such as RACF/SAF). If you use the Service Console on Windows, you must pass the Windows workstation user ID and password prompt. A password is always required to access the Service Console. Even if a user ID is allowed to log into the operating system without a password, access to the Service Console will be denied. If necessary, you must create a password for the user ID that is being used to log in to the Service Console. For more information about the Service Console, see the *IBM Tivoli Monitoring Troubleshooting Guide*.

Communication between components

To secure communication between agents, monitoring servers, and the portal server, use SPIPE as the protocol when you configure communications between the Tivoli Enterprise Portal Server and the hub Tivoli Enterprise Monitoring Server, between hub and remote monitoring servers, and between agents and monitoring servers.

Two additional protocols are used to secure communication between portal clients and the portal server:

- Secure Hypertext Transport Protocol (HTTPS) to retrieve files and Interoperable Object Reference (IOR). The integrated browser in the client provides HTTPS support on the client side; for the server, consider using a third party Web server that supports HTTPS, such as the IBM HTTP Server.
- Internet Inter-ORB Protocol (IIOP) to secure the communications between the portal server and client. This uses Secure Sockets Layer (SSL) provided by MicroFocus. This secure communication uses public key cryptography. See “Using SSL between the portal server and the client” on page 454 for more information.

HTTPS can also be used to secure communication between:

- tacmd Command-Line Interface and the hub Tivoli Enterprise Monitoring Server or portal server.
- Dashboard users and Dashboard Application Services Hub.
- Dashboard Application Services Hub and the dashboard data provider component of the portal server.
- tivcmd Command Line Interface for Authorization Policy and the Dashboard Application Services Hub where the Authorization Policy Server is installed.
- Performance Monitoring service provider and Registry Services, Security Services, and OSLC clients.

If an LDAP server is used with either the hub monitoring server or the portal server, SSL can be used to secure this communication. For more information, see the *Securing Communications* in the *IBM Tivoli Monitoring Administrator's Guide*.

Authorization and authentication

User authentication may be enabled through either the hub Tivoli Enterprise Monitoring Server, or the Tivoli Enterprise Portal Server.

If authentication is enabled through the hub monitoring server, user IDs can be authenticated either by the local operating system registry or by an external LDAP-enabled central registry. User IDs that need to make SOAP Server requests (including user IDs that issue tacmd CLI commands that invoke SOAP server methods) can be authenticated only through the hub monitoring server.

If authentication is enabled through the Tivoli Enterprise Portal Server, user IDs are authenticated against an external LDAP-enabled registry. User IDs that require single sign-on (SSO) capability must be authenticated through the portal server and mapped to unique user identifiers in an LDAP registry shared by all SSO-eligible Tivoli applications.

If you are using Dashboard Application Services Hub with monitoring applications such as IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM or in custom dashboards, you should enable user authentication through the portal server and configure single sign-on if you want to grant your dashboard users different permissions for viewing monitored resources. If you do not configure single sign-on, all dashboard users will see the same set of monitored resources.

The Performance Monitoring service provider component of the Tivoli Enterprise Monitoring Automation Server does not use the user registries configured for the hub monitoring server or portal server to authenticate users. If you want the Performance Monitoring service provider to authenticate HTTP GET requests that it receives from OSLC client applications, you must configure it to use the Security Services component of Jazz for Service Management. Security Services is an optional Jazz for Service Management component that enables non-WebSphere based applications such as the Performance Monitoring service provider to participate in LTPA based single sign-on. See "Single sign-on capability" on page 141 for more details on using Security Services with the Performance Monitoring service provider.

User authentication should not be enabled until at least a basic installation of Tivoli Management Services components and IBM Tivoli Monitoring base agents has been completed and tested. For instructions on enabling authentication, see the *IBM Tivoli Monitoring Administrator's Guide*.

Tivoli Enterprise Portal authorization is controlled by user accounts defined to the portal server. In addition to defining the user IDs that are authorized to log on to the Tivoli Enterprise Portal, these accounts define the permissions that determine the Tivoli Enterprise Portal features a user is authorized to see and use, the monitored applications the user is authorized to see, and the Navigator views (and the highest level within a view) the user can access.

An initial **sysadmin** user ID with full administrator authority is provided during installation so you can log in to the Tivoli Enterprise Portal and add more user accounts. (For information on creating user

accounts and setting user permissions, see the Using Tivoli Enterprise Portal user authorization chapter in the *IBM Tivoli Monitoring Administrator's Guide*.) No password is required to log on to the Tivoli Enterprise Portal, unless user authentication is enabled.

You have two options for authorizing the monitoring resources that can be viewed by users who are using dashboard applications such as IBM Infrastructure Management Dashboards for Servers or custom dashboards:

1. Use the Tivoli Authorization Policy Server and tivcmd Command Line Interface to create roles and permissions, which are collectively called authorization policies. These authorization policies control which managed systems and managed system groups a dashboard user can view. Roles are created for job functions and permissions to view specific managed systems or managed system groups are assigned to roles. Users acquire permissions based on the role (or roles) that the user belongs to. Users can be assigned to roles directly or the user groups that they are members of can be assigned to roles. The permissions also specify the type of object that can be viewed for a managed system or managed system group. The supported object types are event (for situation events) and attribute group (for monitoring data retrieved from an agent).

The following tasks must be performed to use authorization policies:

- a. Install the Tivoli Authorization Policy Server into Dashboard Application Services Hub.
- b. Install the tivcmd Command Line Interface on systems accessible to the administrators who will create and work with authorization policies.
- c. Administrators use the tivcmd Command Line Interface to create authorization policies for dashboard users or user groups.
- d. After the authorization policies have been created for the current set of dashboard users, you must reconfigure the portal server to enable authorization policies. This step causes the portal server to retrieve the authorization policies from the Authorization Policy Server and to start enforcing the authorization policies in the dashboard data provider. If Tivoli Enterprise Portal permissions and monitored application assignments are also configured for the dashboard user, these permissions are ignored since the authorization policies take precedence.

For more information on creating and working with authorization policies, see *Role based authorization policy* in the *IBM Tivoli Monitoring Administrator's Guide*.

2. Use Tivoli Enterprise Portal permissions and monitored application assignments for your dashboard users. If authorization policies are not enabled in the portal server configuration then the dashboard data provider defaults to using Tivoli Enterprise Portal permissions and monitoring application assignments for authorizing dashboard user requests from Dashboard Application Services Hub.

With this option, you create Tivoli Enterprise Portal users for each of your dashboard users using the Tivoli Enterprise Portal User Administration dialog. Using the same dialog, you can grant a user permission to view events and assign the user one of more monitored applications that they can view. These steps can also be performed using the tacmd Command Line Interface. See the *Using Tivoli Enterprise Portal user authorization* chapter in the *IBM Tivoli Monitoring Administrator's Guide* and the *IBM Tivoli Monitoring Command Reference* for more details.

Tivoli Enterprise Portal authorization is less granular than authorization policies. While authorization policies allow you to grant a dashboard user permission to view only specific managed systems or members of specific managed system groups, Tivoli Enterprise Portal authorization is at the monitored application level. In other words, a user is assigned permission to view all managed systems of a particular monitoring application type, for example all Windows OS agents.

When you are initially setting up your monitoring and dashboard environment, it is recommended that you start with Tivoli Enterprise Portal permissions and monitored application assignments. After you are able to see monitoring data in Dashboard Application Services Hub and your administrators have created authorization policies, then reconfigure the portal server if you want to start using authorization policies.

If your dashboard users are also going to access the Tivoli Enterprise Portal client, the set of monitored resources that they can view in dashboards might be different than the monitored

resources they can view in the Tivoli Enterprise Portal client. This can occur if the permissions are inconsistent or the authorization policies are more restrictive.

- Example of inconsistent permissions: Assume the user is granted permission to view a subset of Windows OS agents in Dashboard Application Services Hub using authorization policies but the user is not assigned the Windows OS monitoring application in their Tivoli Enterprise Portal permissions. In this scenario, the user will see their authorized Windows OS agents in the dashboards but they will not see any Windows OS agents when they access the Tivoli Enterprise Portal client.
- Example of more restrictive authorization policies: Assume the user is granted permission to view a subset of Windows OS agents in Dashboard Application Services Hub using authorization policies and the user is assigned the Windows OS monitoring application in their Tivoli Enterprise Portal permissions. In this scenario, the user will see the authorized Windows OS agents in the dashboards but they will see all Windows OS agents when they access the Tivoli Enterprise Portal client.

Dashboard user authorization is also affected by the configuration of the dashboard data provider connection in Dashboard Application Services Hub.

- If the connection is configured for single sign-on, the dashboard users see the monitored resources that they have been authorized to view using either authorization policies, when they are enabled, or Tivoli Enterprise Portal monitored application assignments.
- If the connection is not configured for single sign-on, the Dashboard Application Services Hub always passes the username configured for the connection to the dashboard data provider. Therefore, authorization is performed for the user configured for the connection and not the user who is logged into Dashboard Application Services Hub. In this case, all dashboard users will see the same set of monitored resources.

Because of this behavior, you should configure the dashboard data provider connection for single sign-on if you want to grant different permissions to your dashboard users. Single sign-on is not required if you want all of your dashboard users to have the same authorizations. See “Single sign-on capability” for more information on configuring Dashboard Application Services Hub and the portal server for single sign-on support.

Single sign-on capability

The single sign-on (SSO) feature provides users the ability to launch out of the Tivoli Enterprise Portal to other Tivoli Web-based or Web-enabled applications, or to launch into the Tivoli Enterprise Portal from those applications, without having to re-enter their user IDs and passwords.

Single sign-on is also required if you are using monitoring dashboard applications such as IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM or custom dashboards with Dashboard Application Services Hub and you want to assign different permissions to your dashboard users or launch from the monitoring dashboards to the Tivoli Enterprise Portal client.

For SSO to be enabled for these scenarios, authentication must be configured through the Tivoli Enterprise Portal Server and the LDAP registry defined to the portal server must be a central registry shared by all participating Tivoli applications. All the participating applications must be configured for SSO and must belong to the same internet or intranet domain and realm.

If you are using the Performance Monitoring service provider, it uses the Security Services component of Jazz for Service Management to support single sign-on. When the Performance Monitoring service provider receives an HTTP GET request from an OSLC client, it forwards the LTPA token to Security Services to authenticate the request. If the request does not contain a LTPA token or Security Services indicates that the token is not valid or has expired, the Performance Monitoring service provider returns

an HTTP 401 status code to indicate that the request could not be authenticated. To enable single sign-on support for the Performance Monitoring service provider, perform the following steps:

- Install Security Services on the same application server as the Registry Services component of Jazz for Service Management.
- Enable WebSphere Global Security in the application server for Registry Services and Security Services, configure the application server to use a central LDAP registry and enable single sign-on. For more information, see *Configuring Jazz for Service Management for a central user registry* and *Configuring Jazz for Service Management for SSO* in the Jazz for Service Management Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/psc_ic-homepage.html).
- Configure the application servers for OSLC client applications to use the same LDAP registry and enable LTPA based single sign-on with Security Services. The applications must also be configured to use the same LTPA key.
- Configure the Performance Monitoring service provider to use Security Services by setting the Tivoli Enterprise Monitoring Automation Server KAS_SECURITY_SERVICES_ENABLED environment variable to Yes and restart the automation server.

For instructions on using SSO, see the *Enabling user authentication* chapter in the *IBM Tivoli Monitoring Administrator's Guide*.

SOAP server security

You can control access to the SOAP server in the following two ways.

- You can control *who* is permitted to make requests by enabling user authentication on the hub monitoring server.

If the **Security: Validate User** option is *not* enabled, the SOAP server honors all requests regardless of the sender. If the **Security: Validate User** option on the hub monitoring server is enabled, the SOAP server honors requests only from users defined to the operating system or security authorization facility of the host of the monitoring server.

- You can control *what type of requests* users are permitted to make by configuring the SOAP server.

If you specify a specific type of access for any users, the SOAP server honors requests only from those users, regardless of whether Security: Validate User is enabled.

For information on configuring the security on the SOAP server, see Chapter 17, “Configuring IBM Tivoli Monitoring Web Services (the SOAP Server),” on page 461.

Global Security Toolkit

IBM Tivoli Monitoring includes the Global Security Toolkit (GSKit) for SSL processing as used in SPIPE and HTTPS. GSKit is installed by default on all distributed components, and its utilities are used to create and manage the encryption of data between components through the use of digital certificates.

Note: Do not uninstall or manipulate the GSKit during installation. This may cause functional regression in other products or make them inoperable. The GSKit will automatically install the most recent build if another version GSKit already exists.

On z/OS, GSKit is known as the Integrated Cryptographic Service Facility, or ICSF. If ICSF is not installed and enabled on the z/OS system, the monitoring server uses an alternative, less secure encryption scheme. Since both components must be using the same scheme, if the hub system does not use ICSF, you must configure the Tivoli Enterprise Portal to use the less secure scheme (EGG1) as well. For more information on configuring a monitoring server to use ICSF or to configure the EGG1 scheme, see *Configuring the Tivoli Enterprise Monitoring Server on z/OS* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm). If ICSF is not installed and enabled on a z/OS system for a monitoring server, you must also set the TEMS Security

Compatibility option to Yes (enabled) for that monitoring server so that Take Action requests can be processed by any connected agents that are using the Tivoli Enterprise Monitoring Agent framework at V6.3 or later.

A default certificate and key are provided with GSKit at installation. A stash file provides the database password for unattended operation. You can also use the key management facilities in GSKit to generate your own certificates. For information about the GSKit command line interface, see the GSKCapiCmd Users Guide at ftp://ftp.software.ibm.com/software/webserver/appserv/library/v80/GSK_CapiCmd_UserGuide.pdf.

Note:

1. The IBM Tivoli Monitoring installer no longer modifies the system GSKit. If necessary, it installs a local copy of GSKit that is private to Tivoli Monitoring.
2. In 64-bit environments, only the 64-bit GSKit is installed by default. The 32-bit GSKit is no longer installed by default. The 32-bit GSKit is only installed if a 32-bit component that requires 32-bit GSKit is installed into the 64-bit environment.

Chapter 8. Upgrading from a previous installation

This chapter provides information to help you upgrade from a previous installation.

- “Planning your upgrade” identifies the platforms no longer supported, describes two methods for components need to be upgraded, provides instructions for backing up your existing environment, and specifies the required order for upgrading components.
- “Upgrading from IBM Tivoli Monitoring V6.1 or V6.2” on page 161 provides an overview of the process of upgrading from V6.1 or V6.2 to V6.3 of IBM Tivoli Monitoring.

Planning your upgrade

The following sections provide planning information for upgrading to IBM Tivoli Monitoring V6.2/V6.3.

- “Platforms no longer supported for IBM Tivoli Monitoring V6.3”
- “Prerequisites for IBM Tivoli Monitoring V6.3” on page 146
- “IBM Tivoli Monitoring coexistence and interoperability” on page 158
- “Components to upgrade” on page 149
- “Required order of upgrade” on page 150
- “Migrated information when upgrading from a previous version” on page 152
- “Backing up IBM Tivoli Monitoring” on page 152

Note: In previous versions there was no validation on hub and remote monitoring server protocols and stand-by configuration, so it is possible that your configuration has incorrect entries. Consider reconfiguring all monitoring servers, both hub and remote. During this reconfiguration you will be warned about incorrect entries. Reconfiguration may be particularly useful if you are experiencing long startup and stopping times for monitoring servers. You can reconfigure a hub monitoring server as a remote, but reconfiguring a remote monitoring server from a remote to a hub is currently not supported. In this case you must reinstall the monitoring server.

Platforms no longer supported for IBM Tivoli Monitoring V6.3

The following platforms are no longer supported for IBM Tivoli Monitoring V6.3.

- AIX V5.3 (32 bit): All components
- AIX V5.3 (64 bit): All components
- HP-UX 11i v3 on Integrity (IA64): Tivoli Enterprise Monitoring Server, Warehouse Proxy Agent, Summarization and Pruning Agent
- HP-UX 11i v2 on Integrity (IA64): Tivoli Enterprise Monitoring Server, Warehouse Proxy Agent, Summarization and Pruning Agent
- HP-UX 11i v1 (32/64) on PA-RISC: OS monitoring agent
- Asianux 2.0 on Itanium IA64 (64 bit): All components
- Asianux 3.0 on Itanium IA64 (64 bit): Warehouse Proxy Agent, Summarization and Pruning Agent, Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Portal client, Tivoli Performance Analyzer
- Red Flag 4.1 (32 bit): All components
- Red Flag 5.0 (32 bit): Warehouse Proxy Agent, Summarization and Pruning Agent, Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Portal client, Tivoli Performance Analyzer
- SuSE Linux Enterprise Server 10 on Itanium IA64 (64 bit): OS monitoring agent
- SuSE Linux Enterprise Server 9 Intel x86-32 (32 bit): All components
- SuSE Linux Enterprise Server 9 on Intel x86-64 (64 bit): All components

- SuSE Linux Enterprise Server 9 on Itanium IA64 (64 bit): OS monitoring agent
- SuSE Linux Enterprise Server 9 for iSeries and pSeries (64 bit): OS monitoring agent
- RedHat Enterprise Linux 4 Intel x86-32 (32 bit): All components
- RedHat Enterprise Linux 4 on Intel x86-64 (64 bit): All components
- RedHat Enterprise Linux 4 on Itanium IA64 (64 bit): OS monitoring agent
- RedHat Enterprise Linux 4 on iSeries and pSeries (64 bit): OS monitoring agent
- RedHat Enterprise Linux 4 on zSeries (64 bit): All components
- RedHat Enterprise Linux 4 on zSeries (31 bit): All components
- RedHat Enterprise Linux 5 on Itanium IA64 (64 bit): OS monitoring agent
- Solaris V9 (SPARC) (32/64 bit): All components
- Solaris V9 (SPARC) (64 bit Exploit): OS monitoring agent
- Solaris V8 (SPARC) (64 bit Exploit): OS monitoring agent
- Windows Server 2008 Datacenter Edition on Itanium IA64 (64 bit): OS monitoring agent
- Windows Server 2003 Standard Edition R2 on Intel x86-32 (32 bit): Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Portal client, Summarization and Pruning Agent, Warehouse Proxy Agent, Tivoli Performance Analyzer
- Windows Server 2003 Standard Edition R2 on Itanium IA64 (64 bit): Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Portal client, Summarization and Pruning Agent, Warehouse Proxy Agent, Tivoli Performance Analyzer
- Windows Server 2003 Standard Edition R2 on x86-64 (64 bit): Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Portal client, Summarization and Pruning Agent, Tivoli Performance Analyzer
- Windows Server 2003 Enterprise Edition R2 on Intel x86-32 (32 bit): Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Portal client, Summarization and Pruning Agent, Warehouse Proxy Agent, Tivoli Performance Analyzer
- Windows Server 2003 Enterprise Edition R2 on x86-64 (64 bit): Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Portal client, Summarization and Pruning Agent, Tivoli Performance Analyzer
- Windows Vista Enterprise Edition on Intel x86-32 (32 bit): All components
- Windows Vista Enterprise Edition on Intel x86-64 (64 bit): Portal client
- Windows XP Professional on Intel x86-32 (32 bit): All components
- Windows XP Professional with FDCC on Intel x86-32 (32 bit): All components
- z/OS 1.10 (31/64 bit): All components
- z/OS 1.11 (31/64 bit): All components

Note: When installing on AIX, you can see AIX R5.3 in the messages although it is no longer a supported version. Please ignore the version number; it is a current installer limitation.

For a complete description of the supported platforms for IBM Tivoli Monitoring V6.3, see Chapter 3, “Hardware and software requirements,” on page 75.

Prerequisites for IBM Tivoli Monitoring V6.3

Some of the prerequisites have changed for the current version of IBM Tivoli Monitoring.

Review the Chapter 3, “Hardware and software requirements,” on page 75 for each component and platform and apply all necessary upgrades and fix packs before you begin your upgrade.

Native 64 bit versions of IBM Tivoli Monitoring distributed infrastructure

IBM Tivoli Monitoring V6.3 offers native 64 bit versions of IBM Tivoli Monitoring infrastructure components for Windows, Linux on AMD64 and Intel EMT (x64), Linux on System z, and AIX systems.

On 64 bit systems, it is recommended to install components from the 64 bit version of the installation image. However, you can install components from the 32 bit installation image on 64 bit systems. On 32 bit systems, you can only install components from the 32 bit installation image.

The following conditions apply to upgrading the Tivoli Enterprise Monitoring Server and Tivoli Enterprise Portal Server:

- Automatically upgrading the Tivoli Enterprise Portal Server from 32 bit to 64 bit or from 64 bit to 32 bit is not permitted. For instructions on performing this conversion, see “Upgrading a 32 bit portal server to 64 bit” on page 261 or “Upgrading a 32 bit Windows portal server to a new system running 64 bit Windows” on page 240.
- Upgrading the Tivoli Enterprise Monitoring Server from 64 bit to 32 bit is not permitted.
- Automatically upgrading the Tivoli Enterprise Monitoring Server from 32 bit to 64 bit is permitted.

Note: Certain OMEGAMON agents are designed to run as shared libraries that load into the Tivoli Enterprise Monitoring Server process address space. These agents might not continue to function after you have upgraded the Tivoli Enterprise Monitoring Server from 32 bit to 64 bit on your system. The IBM Tivoli Monitoring installer will recognize this condition on your system and warn you of the possible problem.

If a bit version conflict is detected during the upgrade process, an error message is displayed identifying both the component and the conflict. The Tivoli Enterprise Portal Server, Tivoli Enterprise Monitoring Server, or both, are removed from the list of available features.

Note: If you are upgrading all 32 bit components to IBM Tivoli Monitoring V6.3, and also want to upgrade a 32 bit Tivoli Enterprise Monitoring Server to a 64 bit version, upgrade the components in the following order:

1. Use the 32 bit IBM Tivoli Monitoring installation image to upgrade the Tivoli Enterprise Portal Server first (and 32 bit agents if required).
2. Use the 64 bit IBM Tivoli Monitoring installation image to upgrade the Tivoli Enterprise Monitoring Server to the 64 bit version (and 64 bit agents if required).

Upgrading and Migrating DB2 Database for Linux, UNIX, and Windows

IBM Tivoli Monitoring V6.3 includes a limited-use version of IBM DB2 Enterprise Server Edition 10.1 for use with the Tivoli Enterprise Portal Server and the Tivoli Data Warehouse.

As long as you're using a supported version of DB2 Database for Linux, UNIX, and Windows, upgrading and migrating to IBM DB2 Enterprise Server Edition 10.1 is not required. See “Supported databases for Tivoli Enterprise Portal Server” on page 82 for supported database versions.

If you elect to upgrade and migrate from an earlier version of DB2 for Linux, UNIX, and Windows, you may receive a warning during instance migration that the new edition of DB2 is different from the edition prior to upgrade (Workgroup Edition versus Enterprise Server Edition). As a result, certain DB2 GUI utilities such as db2cc may fail to start after migration, possibly due to the DB2 **jdk_path** parameter being restored to its default value during the migration. To resolve this, **jdk_path** may need to be reset: as the db2inst1 user, run the following DB2 command:

```
db2 update dbm cfg using jdk_path /home/db2inst1/sqllib/java/jdk32
```

See *Configuring DB2 database systems in DB2 Version 10.1 for Linux, UNIX, and Windows* for more information on the **jdk_path** parameter. A complete suite of DB2 product documentation is provided in the IBM DB2 Database information center at <http://pic.dhe.ibm.com/infocenter/db2luw/v10r1/index.jsp>.

Upgrading Tivoli Enterprise Services User Interface Extensions

If you upgrade any OS agent to V6.3 you must also separately upgrade the Tivoli Enterprise Services User Interface Extensions (tacmd CLI) on the same machine to V6.3. If tacmd CLI (product code *ue*) is not upgraded, commands run from this machine will fail.

Tivoli Business Service Manager and Tivoli Enterprise Portal Server integration over SSL

For releases before V.6.2.3, HTTP was only used to transfer static content such as help files from the embedded HTTP server in IBM Tivoli Monitoring, or for single sign-on using the embedded HTTP server in the WAS component. As of version 6.2.3 and higher, the Tivoli Enterprise Portal Server installation contains a required IBM HTTP Server front end for all HTTP and HTTPS traffic.

The IBM HTTP Server provides increased scalability, performance, and unifies HTTP(S) access to the Tivoli Enterprise Portal Server to a single pair of ports.

By default, the IBM HTTP Server is assigned ownership of the HTTP and HTTPS ports 15200 and 15201. These ports are currently used by the embedded HTTP server in the Tivoli Enterprise Portal Server. This ensures that existing Tivoli Enterprise Portal Server clients, configured for single sign-on, migrate to the new architecture. However, with the use of HTTPS you must now consider the SSL/TLS certificate of the server.

A trusted signing authority must sign the IBM Tivoli Monitoring Certificate. You must import the certificate of that signing authority into any web applications that are accessing IBM Tivoli Monitoring data from the Tivoli Enterprise Portal Server using HTTPS. If this certificate is missing, and your Tivoli Enterprise Portal Server clients attempt to connect to port 15201 using SSL/TLS, the error message *no trusted certificate found* is displayed.

The following applications are affected:

- The Tivoli Business Service Manager Charts services for IBM Tivoli Monitoring, that is provided with Tivoli Business Service Manager 4.2.2 and higher.
- IBM Tivoli Monitoring policy-based data fetcher, that is provided with Tivoli Business Service Manager 4.2 and higher.
- Tivoli Integrated Portal when it is configured to use the IBM Tivoli Monitoring charting web service.

Ensuring web applications connecting to the embedded HTTP server in the portal server using SSL continue to work after upgrading to V6.3:

Before you upgrade to V6.3, you must generate a new pair of keys and sign the IBM Tivoli Monitoring Certificate with the key of your signer if you do not want to use the default self-signed certificates installed by IBM Tivoli Monitoring. You then import the public key certificate of your signer into the truststores of your client applications, before the upgrade to V6.3.

Follow these steps to apply this method and minimize the potential for outages:

1. For instructions on how to create a new public/private key pair, requesting a certificate signing, and saving the signed certificate, see the *IBM Tivoli Monitoring Administrator's Guide*.
2. You must update the truststore of the client with the signer's certificate, if this has not already been done. The instructions in this section are specific to the Tivoli Business Service Manager Charts services for IBM Tivoli Monitoring that is provided with Tivoli Business Service Manager 4.2, 4.2.1 Fix Pack 2.

For the Tivoli Business Service Manager data server:

1. Log in to the server where the Tivoli Business Service Manager data server is installed.
2. Change directory to `TIP_HOME/bin`.

3. Issue the following command:
`./wsadmin.sh ?profileName TBSMProfile ?user tipadmin`
4. When prompted, enter the password for the tipadmin user.
5. From the wsadmin prompt, issue the following command:
`$AdminTask retrieveSignerFromPort { -host tepshostname -port 15201 -keyStoreName NodeDefaultTrustStore -certificateAlias alias-name }`

Where:

tepshostname

Is the hostname of the Tivoli Enterprise Portal Server.

alias-name

Is an alias for the certificate in the NodeDefaultTrustStore, for example *TEP-IHS*.

6. Save the configuration:
`wsadmin>$AdminConfig save`
7. Exit the wsadmin prompt:
`wsadmin>quit`
8. Stop and restart the Tivoli Business Service Manager data server for the changes to take effect.

For the Tivoli Business Service Manager Dashboard server or a Tivoli Integrated Portal server that uses the IBM Tivoli Monitoring charting web service:

1. Log in to the WebSphere Application Server administrative console.
2. Expand **Security** and click **SSL certificate and key management**.
3. Click **Keystores and certificates**.
4. Click **NodeDefaultTrustStore**.
5. Click **Signer certificates**.
6. Click **Retrieve from port**.
7. Enter the host name of the Tivoli Enterprise Portal Server.
8. Enter port **15201**.
9. Enter an alias name, for example *TEP-IHS*.
10. Click **Retrieve signer information**.
11. Click **OK**.
12. Stop and restart the Tivoli Business Service Manager dashboard server or Tivoli Integrated Portal server for the changes to take effect.

Components to upgrade

You can determine what components to upgrade by using either of two methods.

About this task

Use the Enterprise level of the Managed Systems Status list to display the version of any connected agent, as well as its status (online/offline). Identifying an agent's status is especially important as you plan remote deployments. An alternative method, available at the Enterprise level of the Navigator, is to use the Tivoli Management Services Infrastructure view. This topology view visually expresses the relationships and linking of monitoring agents and other components to the hub monitoring server. Use hover-help (help flyovers) in the Tivoli Enterprise Portal topology view to determine the current version of your installed monitoring agents and other components.

Procedure

1. Access the Tivoli Enterprise Portal through the desktop client or the browser client:
 - **Desktop client:**
 - a. In the Windows **Start** menu, select **Programs → IBM Tivoli Monitoring → Tivoli Enterprise Portal**. The login message is displayed.
 - b. Type the name (sysadmin) and password of the account that you created during installation.
 - c. Click **OK**.
 - **Browser client:**
 - a. Start the browser.
 - b. Type the following URL for the Tivoli Enterprise Portal into the Address field of the browser:
`http://systemname:15200/`
where the *systemname* is the host name of the computer where the Tivoli Enterprise Portal Server is installed, and 15200 is the port number for the browser client. 15200 is the default port number for the browser client. Your portal server might have a different port number assigned.
 - c. Click **Yes** in response to any security prompts.
 - d. In the login prompt, type the user name and password of the sysadmin account that you created during installation.
 - e. Click **OK**.
 - f. Click **Yes** to accept Java Security certificates for this browser session.
2. Click the Navigator, the view that is located by default in the upper-left corner of the portal. The first display that you see in the Navigator view is Enterprise Status workspace.
3. Select the **Enterprise** node at the top of the navigation tree, if it is not already selected.
4. Right-click the node and select **Workspace → Self-Monitoring Topology** to display the default topology view.
5. Click on the **TMS Infrastructure - Base view**. This view contains the topology display.
6. Click on the Maximize icon in the upper-right corner of the portal. The topology view occupies the entire display area of the workspace.

Required order of upgrade

Upgrade the products that comprise your configuration in the order described in this section.

If multiple components are being installed within a single installation directory, upgrade all components at once rather than running the upgrade multiple times.

1. Event synchronization component

If you use event synchronization, upgrade the event synchronization component first. If you are using Netcool/OMNIbus, refer to Chapter 30, "Setting up event forwarding to Netcool/OMNIbus," on page 697. If you are using Tivoli Enterprise Console, refer to Chapter 31, "Setting up event forwarding to Tivoli Enterprise Console," on page 799.
2. Tivoli Data Warehouse
 - a. Stop all instances of the Warehouse Proxy Agent and the Summarization and Pruning Agent.
 - b. If you are upgrading from V6.1.x and are collecting warehouse data from any of the agents listed in Table 27 on page 151, you must migrate your warehouse database before you upgrade other components. See "Upgrading the warehouse from IBM Tivoli Monitoring V6.1.x levels" on page 156 and the user's guide for each agent for details.

Table 27. Agents requiring warehouse database migration

- Windows OS agent
- Unix OS agent
- Monitoring for Databases: DB2 Database for Linux, UNIX, and Windows agent
- Monitoring for Databases: Sybase agent
- Monitoring for Databases: mySAP agent
- Monitoring for Databases: Oracle agent

Install the Tivoli Enterprise Portal Server (step 5) before you upgrade the warehouse components to ensure the SQL scripts needed for the warehouse upgrade are available. Ensure you upgrade *all* your site's Warehouse Proxy Agents and Summarization and Pruning Agents before performing step 8.

Note: Running the upgrade scripts a second time may cause loss of warehouse data.

If you are upgrading from V6.2.x, you do not need to perform these migration steps.

3. Hub Tivoli Enterprise Monitoring Server

Upgrading the hub monitoring server is not required, but if the hub is not upgraded, new features of upgraded agents will not be available. If the hub is upgraded, it must be upgraded first. If the Hot Standby feature is used, the primary and secondary hub monitoring servers must be upgraded in that order before the remote monitoring servers are upgraded.

Note: Installing a back-level OS agent into an existing V6.3 environment is not supported. If this is the desired configuration, install the back-level OS agent first.

4. Tivoli Enterprise Monitoring Automation Server

5. Remote monitoring servers (if necessary)

6. Tivoli Enterprise Portal Server

7. Tivoli Enterprise Portal desktop client

The client and the portal server must be at the same level for the client to connect successfully.

8. Infrastructure Management Dashboards for Servers

9. Tivoli Authorization Policy Server

10. tivcmd Command Line Interface for Authorization Policy

11. Tivoli Performance Analyzer:

- Tivoli Performance Analyzer must be at level V6.2.2 fix pack 2 or later before installing V6.3.
- For configuration information, see step 16 on page 178 in "Installation procedure" on page 172.
- For information about historical data collection settings when upgrading from V6.2.2, see note 4 in "Installing domain definitions for Tivoli Performance Analyzer" on page 336.

12. Monitoring agents

Note: The appropriate level of the Tivoli Enterprise Management Agent Framework is installed when an agent is installed.

Self-describing monitoring agents is a feature that integrates the installation of an agent with the dispersal and installation of associated product support files throughout your IBM Tivoli Monitoring infrastructure. To use the self-describing monitoring function, the Tivoli Enterprise Management Agent Framework must be at version 6.2.3 or later. For more information, see "Self-describing agent installation" on page 394.

Note:

1. Agents created with the Agent Builder require that the entire path from the agent, through the remote monitoring server it connects to, to the hub monitoring server, to the portal server, and the data warehouse be upgraded to V6.2 or later.

2. Installing components or agents in an existing CANDLEHOME or installation directory is supported as long as the user ID used to run the installation is always the same. Installing components or agents in an existing CANDLEHOME or installation directory using different user IDs is not supported.

Migrated information when upgrading from a previous version

When installing over a previous release (into the same installation directory) the following information is migrated into the new version.

- **Windows**
 - Port number and communication protocol settings
 - Situations
- **UNIX** Situations

Backing up IBM Tivoli Monitoring

Before upgrading, make a backup copy of your current IBM Tivoli Monitoring installation in case you need to revert to it.

You should always contact IBM Software Support before attempting to use the files generated by the following procedures to restore the IBM Tivoli Monitoring environment. The support staff can help ensure success of the restoration. Otherwise, errors made during the modification of the Windows registry could lead to a corrupted operating system or other problems.

The following procedures describe valid methods for creating a backup. However, an enterprise typically has a standard process for preserving a backup image of all computer systems, including the IBM Tivoli Monitoring environment. In most cases, the standard process for your enterprise is preferable to the processes described in this section. Use the instructions in this section only when a standard backup process is not available.

You must perform the following instructions when no other activity is running on the system. The user accounts that you use to perform these activities must have **root** or **Administrator** privileges:

- “Backing up a Windows installation”
- “Backing up a UNIX or Linux installation” on page 154

Backing up a Windows installation

This section contains instructions for backing up your Windows installation.

Before you begin

Typically an enterprise has a standard process for preserving a backup image of all computer systems, including the IBM Tivoli Monitoring environment. In most cases, the standard process for your enterprise is preferable to the processes described in this section. Use the instructions in this section only when a standard backup process is not available.

About this task

Follow these steps to back up a Windows installation. In step 4 on page 153 of this procedure, you record the current configuration settings. In Step 9 on page 153, you apply these settings again.

Procedure

1. Close any browser or desktop clients for the Tivoli Enterprise Portal.
2. Launch the **Manage Tivoli Enterprise Monitoring Services** (KinConfig.exe) utility.
3. On the computer where you are backing up an installation, stop the Tivoli Enterprise Portal Server, the Tivoli Enterprise Monitoring Server, the Eclipse Help Server, and all the monitoring agents running on the system.

4. Perform the following steps to record the current configuration settings for IBM Tivoli Monitoring. (For example, you can make a record on a sheet of paper or in a buffer file.)
 - a. Right-click in the row of the component that you want to configure.
 - b. Select **Reconfigure**.

Note: You make no changes in the series of windows that are displayed.

- c. Record the settings for the component in the series of configuration windows.
You record details such as port numbers, host names, protocols for communication, firewall settings, and settings for the data warehouse.
- d. Click **OK** in each window and accept all prompts, *without making changes*.

Note: If you click **Cancel** at any point, the **Reconfigure** process fails to display all configuration windows. You must start the **Reconfigure** process again.

You must *unconfigure* the monitoring environment in the next step. This action restores the Windows registry to a known state that ensures success, if you restore the environment later.

5. Perform the following steps to unconfigure each IBM Tivoli Monitoring component, except the Tivoli Enterprise Portal desktop client. Do not perform the following steps for the desktop client.
 - a. Right-click in the row of a component in the **Manage Tivoli Enterprise Monitoring Services** window.
 - b. Select **Advanced >> Unconfigure** in the popup menu.

When you are finished, all components except the Tivoli Enterprise Portal desktop show **No** in **Configured** column of the **Manage Tivoli Enterprise Monitoring Services** window.

6. Use a compression command to compress the contents of the directory where IBM Tivoli Monitoring is installed.
7. Use the appropriate database commands to back up the Tivoli Enterprise Portal Server and Tivoli Data Warehouse databases.
For more information, see "Backing up your portal server and Tivoli Data Warehouse databases" on page 155.
8. Export the entire Windows registry to a backup file as follows:
 - a. Select **Run** in the Windows **Start** menu.
 - b. Type **regedit** in the **Open** field.
 - c. Click **OK**. The Registry Editor is displayed.
 - d. Ensure that the **MyComputer** node at the top of the registry is selected, so that all values in the registry are exported.
 - e. Select **Export** in the **File** menu.
 - f. Save the copy of the registry to the following path and filename: C:\WinRegistryBeforeInstall.reg

At this time, the backup process is complete.

9. Perform the following steps to return IBM Tivoli Monitoring to its normal configuration and status:
 - a. Right-click in the row of each unconfigured component.
"Unconfigured" components show **No** in the **Configured** column of the **Manage Tivoli Enterprise Monitoring Services** window.
 - b. Select **Advanced >> Configure Advanced** in the popup menu.
 - c. In the series of dialog boxes that is displayed, enter the original parameter settings that you recorded in Step 4.
 - d. Click **OK** in each configuration window and accept other prompts to save your changes.

When you are finished, all components show **Yes** in **Configured** column and **Stopped** in the **Status** column.

Results

You are ready to proceed with the upgrade on this computer.

What to do next

Always contact IBM Software Support before attempting to use the files generated in this procedure to restore the IBM Tivoli Monitoring environment. The support staff can help ensure success of the restoration. Otherwise, errors made during the modification of the Windows registry could lead to a corrupted operating system.

Backing up a UNIX or Linux installation

Follow the steps in this section to back up a UNIX or Linux installation.

Before you begin

Normally an enterprise has a standard process for preserving a backup image of all computer systems, including the IBM Tivoli Monitoring environment. In most cases, the standard process for your enterprise is preferable to the processes described in this section. Use the instructions in this section only when a standard backup process is not available.

Procedure

1. Close the Tivoli Enterprise Portal browser and desktop clients.
2. Stop the Tivoli Enterprise Portal Server, the Tivoli Enterprise Monitoring Server, the Eclipse Help Server, and all the monitoring agents running on the system.
3. If the Tivoli Enterprise Portal Server is installed, run the following command:

```
./itmcmd execute cq "runscript.sh migrate-export.sh"
```

Note:

- a. If you are upgrading from IBM Tivoli Monitoring V6.2.2 FP2-FP8, V6.2.3, or V6.2.3 FP1 to V6.3 or later and you have Tivoli Data Warehouse summarization and pruning configured, you should follow the instructions in the Local fix section of the APAR IV12595 before executing the migration-export script. For more information, go to <http://www-01.ibm.com/support/docview.wss?uid=swg1IV12595>.
 - b. As an alternative to using the migrate-export script to backup a DB2 portal server database, you can use the DB2 commands specified in "Backing up your portal server database" on page 155.
4. Use the tar command to compress the contents of *CANDLE_HOME* (the directory where IBM Tivoli Monitoring is installed), using a command that is similar to the following:

```
tar -cvf /tmp/CANDLE_HOME.backup.tar CANDLE_HOME
```

5. Add the following files to the tar file created in step 4 above:

- On AIX:

```
/etc/rc.itm*  
tar -uvf /tmp/CANDLE_HOME.backup.tar /etc/rc.itm*
```

- On HP-UX:

```
/sbin/init.d/ITMAgents*  
tar -uvf /tmp/ITMinstall_dir.backup.tar /etc/init.d/ITMAgents*
```

- On other UNIX or Linux systems:

```
/etc/init.d/ITMAgents*  
tar -uvf /tmp/CANDLE_HOME.backup.tar /etc/init.d/ITMAgents*
```

6. Use the appropriate database commands to back up the Tivoli Data Warehouse databases.

For more information, see "Backing up your portal server and Tivoli Data Warehouse databases" on page 155.

Results

You are ready to proceed with the upgrade on this computer.

What to do next

Always contact IBM Software Support before attempting to use the files generated in this procedure to restore the IBM Tivoli Monitoring environment. The support staff can help ensure success of the restoration. Otherwise, errors made during the modification of the Windows registry could lead to a corrupted operating system.

Backing up your portal server and Tivoli Data Warehouse databases

The following sections contain procedures for backing up your portal server database and backing up a DB2 for Linux, UNIX, and Windows Tivoli Data Warehouse database. If you use Microsoft SQL Server or Oracle for your Tivoli Data Warehouse database, use the corresponding commands.

- “Backing up your portal server database”
- “Backing up your Tivoli Data Warehouse database” on page 156

Note: These steps are required if you are updating a Tivoli Monitoring V6.1 environment to a later release. These steps are recommended (but not required) if you are updating a Tivoli Monitoring V6.2 or later environment to a later release.

Backing up your portal server database

Follow the appropriate procedure in this section to back up your portal server database if you have not already performed the backup described in the procedures in “Backing up IBM Tivoli Monitoring” on page 152.

Use the following procedure to backup your portal server database:

1. Stop Tivoli Enterprise Portal Server
2. Export the Tivoli Monitoring data in the database by using the migrate-export script, as described in the *IBM Tivoli Monitoring Administrator's Guide* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/index.jsp?topic=%2Fcom.ibm.itm.doc_6.3%2Fadminuse%2Ftpsreplicate_export.htm).

Note: If you are upgrading from IBM Tivoli Monitoring V6.2.2 FP2-FP8, V6.2.3, or V6.2.3 FP1 to V6.3 or later and you have Tivoli Data Warehouse summarization and pruning configured, you should follow the instructions in the Local fix section of the APAR IV12595 before executing the migration-export script. For more information, go to <http://www-01.ibm.com/support/docview.wss?uid=swg1IV12595>.

Alternatively you can follow the database specific methods that are described in the sections that follow to back up the portal server database.

DB2 backup procedure

Use the following command to backup your portal server database where DB2 Database for Linux, UNIX, and Windows is the relational database management system of choice.

```
db2 backup database yourtepsdatabase to /yourbackuplocation
```

If an existing connection prevents you from backing up the database, use the following commands.

1. db2 connect to *yourtepsdatabase*
2. db2 quiesce database immediate force connections
3. db2 connect reset
4. db2 backup database *yourtepsdatabase* to */yourbackuplocation*

5. db2 connect to *yourtepsdatabase*
6. db2 unquiesce database
7. db2 connect reset

Derby backup procedure

Use the following procedure to back up your portal server database where the embedded database, Derby, is the relational database management system of choice.

1. Stop the Tivoli Enterprise Portal Server.
2. To manually back up the entire Derby database, copy the database directory to a backup location. This database is stored in either of these locations:

- **Windows** `ITMHOME\CNPSJ\derby\TEPS0`
- **Linux** **UNIX** `ITMHOME/platform/iw/derby/TEPS0`

Backing up your Tivoli Data Warehouse database

Follow the procedure in this section to back up your Tivoli Data Warehouse database.

Use the following command to back up your Tivoli Data Warehouse database where DB2 Database for Linux, UNIX, and Windows is the relational database management system of choice.

```
db2 backup database yourwarehousedatabase to /yourbackuplocation
```

If an existing connection prevents you from backing up the database, use the following commands.

1. db2 connect to *yourwarehousedatabase*
2. db2 quiesce database immediate force connections
3. db2 connect reset
4. db2 backup database *yourwarehousedatabase* to */yourbackuplocation*
5. db2 connect to *yourwarehousedatabase*
6. db2 unquiesce database
7. db2 connect reset

Sites using the Summarization and Pruning Agent with DB2 for Linux, UNIX, and Windows

Your migrated Tivoli Data Warehouse database requires subsequent checking to ensure it is set up so the Summarization and Pruning Agent can perform multiple system batching. Complete these steps:

1. Back up the database, if necessary.
2. Edit the KSY_DB2_WAREHOUSEMARKER.sql script. If using archive logging, modify it as suggested in the script to avoid the need for extra backup at the end.
3. Execute the script:
`db2 -tvf KSY_DB2_WAREHOUSEMARKER.sql`

The code will detect if the table is correctly migrated. If not, a single managed system will be enforced independent of the setting to prevent database deadlocks.

This procedure affects only the WAREHOUSEMARKER table. If you do not complete it, the Summarization and Pruning Agent will do single system batching only.

Upgrading the warehouse from IBM Tivoli Monitoring V6.1.x levels

Note: Complete these steps only if you are upgrading from IBM Tivoli Monitoring V6.1.x levels. They do not apply to any level after and including V6.2.0. Running the upgrade scripts again may lead to loss of some data.

Some of the monitoring agents have made changes to the warehouse tables. There are three types of changes, two of these types of changes require performing upgrade procedures before running the V6.3 Warehouse Proxy and Summarization and Pruning Agents. The procedure for the other change can be performed before or after you upgrade IBM Tivoli Monitoring. These procedures are accomplished using product-provided scripts.

Case 1 changes affect the table structure, and Case 3 changes affect the table location. Both Case 1 and Case 3 changes must be completed before running the 6.3 Warehouse Proxy and Summarization and Pruning Agents. These changes and procedures are documented comprehensively in the following user's guide appendix for each monitoring agent that is affected: "Upgrade: upgrading for warehouse summarization."

- Case 1 changes add a new column to a raw table, and assign a default value to the column. If the 6.3 Warehouse Proxy Agent is started before these scripts are run, a NULL default value is assigned, and that default value cannot be changed. The following monitoring agents have Case 1 changes:
 - Monitoring for Databases: DB2 for Linux, UNIX, and Windows Agent
 - Monitoring for Applications: mySAP Agent
 - Monitoring: UNIX OS Agent
 - Monitoring: Windows OS Agent
- Case 2 changes affect how agent data is summarized because of changes in primary key definitions. The Case 2 changes can be done before or after you upgrade IBM Tivoli Monitoring. The following monitoring agents have Case 2 changes:
 - Monitoring for Applications: mySAP Agent
 - Monitoring for Databases: Sybase Server Agent
 - Monitoring: Windows OS Agent
- Case 3 changes are only applicable if you are using DB2 for Linux, UNIX, and Windows as your warehouse database. These changes move a summary table from the 4K table space to the 8K table space. If the 6.3 Summarization and Pruning Agent is started before these scripts are run, you receive DB2 for Linux, UNIX, and Windows errors, because the row size of the upgraded table is longer than what can fit into the 4K table space. The following monitoring agents have Case 3 changes:
 - IBM Tivoli Monitoring for Databases: DB2 for Linux, UNIX, and Windows Agent
 - IBM Tivoli Monitoring for Databases: Oracle Agent
 - IBM Tivoli Monitoring: UNIX OS Agent

To ensure that you can make the table upgrades required by Case 1 and Case 3 before installing the 6.3 monitoring agents, do the following:

1. Stop the Warehouse Proxy and the Summarization and Pruning Agents prior to upgrading IBM Tivoli Monitoring.
2. Change the default startup of these monitoring agents to "Manual" so they do not start automatically during any restart of host computer. This prevents these agents from starting automatically during the upgrade process.
3. Perform the IBM Tivoli Monitoring upgrades.
4. Make the warehouse updates for the monitoring agents with Case 1 and Case 3 changes as described in the user's guide for the affected agent.
5. Re-enable the automatic startup of the Warehouse Proxy and the Summarization and Pruning Agents, and start these agents.

Several monitoring agents made changes to the warehouse collection and summarization characteristics for some agent attribute groups. These changes correct and improve the way warehouse data is summarized, producing more meaningful historical reports. For an explanation of these changes and the implications to your warehouse collection and reporting, see the "Upgrading for warehouse summarization" appendix of the user's guide for your specific monitoring agent.

IBM Tivoli Monitoring coexistence and interoperability

The statements in this section specify the level of interoperability for IBM Tivoli Monitoring V6.3 in support of existing agents and staged upgrade scenarios.

Special interoperability notes for Tivoli Monitoring V6.2 are provided where appropriate.

Note: If you are upgrading from V6.1 or V6.2.x to V6.3, you must update both infrastructure components and OS agents residing in the same \$ITM_HOME installation at the same time.

Tivoli Enterprise Monitoring Server

Coexistence and interoperability statements for the Tivoli Enterprise Monitoring Server.

- A V6.3 hub monitoring server supports connections from a V6.1, V6.2, V6.2.1, V6.2.2, or V6.2.3 remote monitoring server.
- If the Hot Standby feature is used, the primary and secondary hub monitoring servers must be at the same level.
- The recommended configuration is for the remote monitoring servers to connect to a hub monitoring server at the same or later release level. V6.3 remote monitoring servers can connect to a V6.2, V6.2.1, V6.2.2, or V6.3 (recommended) hub monitoring server. Remote monitoring servers will automatically switch into compatibility mode to match the feature set of an older release level hub monitoring server.

Note: Agents created with the Agent Builder require that the entire path from the agent, through the remote monitoring server to which it connects, to the hub monitoring server, to the portal server and the data warehouse be at IBM Tivoli Monitoring V6.2 or later.

Also, agents built with the V6.2.1.x Agent Builder (or previous) cannot be installed into an environment where a System Monitor Agent is also running.

- If your site uses IP.SPIPE for Tivoli Enterprise Monitoring Server communication:
 - A V6.2.2 or later remote monitoring server not running in FIPS mode can connect to V6.2.1 or earlier hub monitoring server. Connectivity will default to the Triple DES standard. To enable AES encryption on the hub, customers may set parameter GSK_V3_CIPHER_SPECS="352F0A" on the hub monitoring system.
 - A V6.2.1 or earlier remote monitoring server can connect to a V6.2.2 or later hub not running in FIPS mode. Connectivity will default to the Triple DES standard. To enable AES encryption, customers may set parameter GSK_V3_CIPHER_SPECS="352F0A" on the remote monitoring systems.
 - Agents using pre-V6.2.2 monitoring agent framework can select IP.SPIPE to connect to a V6.2.2 or later monitoring server defined with FIPS=N.
 - Agents using pre-V6.2.2 monitoring agent framework cannot select IP.SPIPE to connect to a V6.2.2 or later monitoring server defined with FIPS=Y.
 - Agents using V6.2.2 or later monitoring agent framework defined with either FIPS=N or FIPS=Y can select IP.SPIPE to connect to a V6.2.2 or later monitoring server defined with either FIPS=N or FIPS=Y.
- If you are running a hub monitoring server at release 6.2 or earlier, the tacmd CLI commands for situation groups, situation long names, and adaptive monitoring cannot be used.
- If you are running a hub monitoring server at a release prior to V6.2.3, you cannot use the self-describing agent function. If you are running a remote monitoring server at a release prior to V6.2.3, you cannot use the self-describing agent function for the agents connected to that monitoring server.
- If you are running a remote monitoring server at a release prior to 6.2.2 fix pack 2, you cannot use the tacmd CLI commands getFile, putFile, and executeCommand.
- If you are running a remote monitoring server at a release prior to V6.2.3, you cannot use the tacmd CLI commands for prerequisite checking and setAgentConnection.

- If you are running a remote monitoring server at a release prior to V6.3, installation for the agent(s) that connect to that monitoring server while the self-describing agent installation was suspended (via `tacmd suspendsda`) will not be restarted when processing is resumed (via `tacmd resumesda`). You must recycle the agent or the remote monitoring server.
- If you are running a hub monitoring server at a release prior to V6.3, the auditing of take actions will not be enabled.

Tivoli Enterprise Portal Server

Coexistence and interoperability statements for the Tivoli Enterprise Portal Server.

- The V6.3 Tivoli Enterprise Portal Server can interoperate with older hub Tivoli Enterprise Monitoring Servers back to V6.1. However, it is recommended best practice that the portal server match the version of the hub Tivoli Enterprise Monitoring Server.
- The V6.3 Tivoli Enterprise Portal Server supports an upgrade-in-place from IBM Tivoli Monitoring V6.2, V6.2.1, V6.2.2, or V6.2.3 portal server on the same platform.
- The V6.3 Tivoli Enterprise Portal Server supports an upgrade from IBM Tivoli Monitoring V6.1 portal server on the same platform by importing an IBM Tivoli Monitoring V6.1 database and then upgrading it by running `buildpresentation.bat` (for Windows platforms) or `InstallPresentation.sh` (for Linux/UNIX platforms).
- Migrating a Tivoli Enterprise Portal server to a new platform, including switching from a 32-bit installation to a 64-bit installation on the same operating system, is accomplished with database export and import. For information about migrating the database, see the description of the `migrate-export` and `migrate-import` scripts in the *IBM Tivoli Monitoring Administrator's Guide*.
- When upgrading from a prior portal server to the current version, it is important that your system's temporary directory has sufficient space to accommodate the download of the eWAS server code and the Eclipse server code. See the memory and disk space requirements documented in Table 18 on page 85.
- The Tivoli Enterprise Portal clients must be at the same release level as the portal server.
- If you are using MS SQL Server as your portal server database, the portal server database schema name must match the portal server user ID. Otherwise the Tivoli Enterprise Portal Server will fail to start and an error message is displayed. This is not an issue for a pristine installation of IBM Tivoli Monitoring. However, if you are upgrading your environment from V6.2.2 or earlier, it is possible that the database schema name and user ID do not match. If this is the case, change the portal server database schema name to match the portal server user ID.

Tivoli Data Warehouse

Coexistence and interoperability statements for the Tivoli Data Warehouse.

- All Warehouse Proxy and Summarization and Pruning Agents must be at the same level to avoid inconsistencies in presentation.
- Warehouse Proxy and Summarization and Pruning agents can interoperate with monitoring servers at the same or later release level.
- It is recommended that you upgrade the Warehouse Proxy and Summarization and Pruning agents to be at the same release level as the most recent agent monitoring framework release level used by your agents. For example, if you intend to use IBM Tivoli Monitoring V6.2.3 operating system agents, then upgrade the Warehouse Proxy and Summarization and Pruning agents to V6.2.3 as well. This ensures the warehouse agents support all of the history collection functions used by your agents:
 - The V6.2.1 or later Warehouse Proxy Agent and Summarization and Pruning Agent are required for correct functioning of any agent requiring 64-bit integer data support.
 - The V6.2.2 or later Warehouse Proxy Agent and Summarization and Pruning Agent are required if you configure granular history collection for any of your agents using the V6.2.2 or later agent monitoring framework.

- The V6.2.3 or later Warehouse Proxy Agent and Summarization and Pruning Agent are required if you configure historical data compression before upload to the warehouse database for any of your agents using V6.2.3 or later agent monitoring framework.

Agents

Coexistence and interoperability statements for Agents.

IBM Tivoli Monitoring V6.3 agents (and application agents using the V6.3 agent monitoring framework) require a V6.3 Tivoli Enterprise Monitoring Server and Tivoli Enterprise Portal Server. In other words, the V6.2.3 (and earlier) monitoring server and portal server do not support V6.3 agents.

Pre-V6.3 agents cannot coexist on the same machine as V6.3 Tivoli Monitoring server components, and must be upgraded at the same time as the V6.3 Tivoli Monitoring server components.

Back-level agents can connect through up-level monitoring servers. In particular, V6.2.1 agents will work with the V6.2.1, V6.2.2, V6.2.3, and V6.3 Tivoli Enterprise Monitoring Server and Tivoli Enterprise Portal Server.

Agents that use dynamic affinity cannot connect or fail over to a pre-V6.2.1 remote monitoring server.

An agent upgrade elevates the user credentials of agents that have been set at deployment time to run with reduced (non-root) credentials. Use one of the following steps to stop the agent and restart it with the appropriate user credentials:

- After stopping the agent, log in (or invoke su) as the desired user, and then run the itmcmd command to start the agent in that user's context.
- Edit the startup file, and add the following line to change the default user context for starting agents:

```
/usr/bin/su - dbinstancename -c "itmhome/bin/itmcmd agent -h itmhome  
-o dbinstancename start product_code
```

where:

product_code

is the two-character product code for the agent (for example, ud for DB2 for Linux, UNIX, and Windows).

The Tivoli Monitoring Universal Agent and Tivoli Monitoring UNIX Log Alert Agent are no longer provided with IBM Tivoli Monitoring V6.3. The V6.1.x and V6.2.x versions of those agents will continue to connect to V6.3 monitoring servers.

Tivoli Event Synchronization component

Coexistence and interoperability statements for the Tivoli Event Synchronization component.

If you have already installed a previous version of Tivoli Event Synchronization, you must upgrade to the latest version. If you are using Netcool/OMNIbus, refer to “Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIbus integration” on page 766. If you are using Tivoli Enterprise Console, refer to “Upgrading to Tivoli Event Synchronization version 3.0.0.0” on page 826. The upgrade is required for the event server to correctly parse events coming from the V6.3 hub monitoring server and from monitoring agents configured to send private situation events.

IBM Tivoli Monitoring V5.x interoperability

Coexistence and interoperability statements for V5.x.

Using the IBM Tivoli Monitoring 5.x Endpoint agent, you can view data from IBM Tivoli Monitoring 5.x resource models in the Tivoli Enterprise Portal and warehouse granular data in the Tivoli Data

Warehouse. You can use this visualization to replace the Web Health Console used in IBM Tivoli Monitoring V5.1. For information about using this visualization, see the *IBM Tivoli Monitoring Agent for IBM Tivoli Monitoring 5.x Endpoint User's Guide*.

Upgrading from IBM Tivoli Monitoring V6.1 or V6.2

The following sections provide information for upgrading from IBM Tivoli Monitoring V6.1 or V6.2 to V6.3.

- “Overview of the upgrade process”
- “Required Java Runtime Environment” on page 165

Overview of the upgrade process

This section outlines the steps for upgrading from IBM Tivoli Monitoring V6.1 or V6.2 to IBM Tivoli Monitoring V6.3.

Table 28. Upgrading from IBM Tivoli Monitoring V6.1 or V6.2 to IBM Tivoli Monitoring V6.3

	Task	Where to find information
Before upgrading your monitoring environment	Review the upgrade planning information to identify what you can and cannot upgrade, as well as the required upgrade order.	“Planning your upgrade” on page 145
	Linux and UNIX sites: If your Tivoli Enterprise Portal Server is running as a non-root process and you plan to upgrade it using a non-root userid, then a special procedure must be performed prior to upgrade.	“Linux and UNIX: Upgrading a portal server running as a non-root process” on page 163
	Stop all components that you are upgrading and change their startup from Automatic to Manual . Stop Manage Tivoli Enterprise Monitoring Services, and do not start any other Tivoli Management Services components while you complete this upgrade.	“Starting and stopping components” on page 420
	Restart the computer on which you are installing IBM Tivoli Monitoring.	
Upgrading your monitoring environment	Run the IBM Tivoli Monitoring installation program on all components that you want to upgrade. Use your existing installation directory as your IBM Tivoli Monitoring directory.	Chapter 10, “Installing IBM Tivoli Monitoring,” on page 193

Table 28. Upgrading from IBM Tivoli Monitoring V6.1 or V6.2 to IBM Tivoli Monitoring V6.3 (continued)

	Task	Where to find information
After upgrading your monitoring environment	On platforms other than Windows, you must reconfigure the Tivoli Enterprise Monitoring Server and the Tivoli Enterprise Portal Server. ² Note: On Windows platforms, the installer manages the reconfiguration process.	Chapter 13, “Configuring IBM Tivoli Monitoring components,” on page 411
	On Windows platforms, if you chose upgrade seeding, support for base agents is upgraded automatically, but you must add application support for any other monitoring agents you are upgrading.	“Configuring application support for nonbase monitoring agents” on page 279
	<p>After upgrading your Tivoli Enterprise Monitoring Server, you must add application support to it.² This step updates existing situation definitions and installs new situations provided with Tivoli Monitoring V6.3.</p> <p>Application support is contained in the ms component.</p> <ul style="list-style-type: none"> • Windows You are automatically prompted to add application support to the monitoring server using the kms_upg.sql file. • Linux UNIX You must initiate upgrade reseeding using the itmcmd support command and specifying the ms component: <pre>./itmcmd support -t tems_name ms</pre> <p>When reseeding completes, recycle your monitoring server:</p> <pre>./itmcmd server stop tems_name ./itmcmd server start tems_name</pre> 	“Installing and enabling application support” on page 275
<p>Note:</p> <ol style="list-style-type: none"> 1. The installation instructions sometimes reference the default path for the installation of IBM Tivoli Monitoring. If you did not install the product in the default path, you must specify the correct path whenever the upgrade procedures require a path specification. 2. After you upgrade a UNIX monitoring server, you must run the itmcmd config -S command to configure the upgraded monitoring server. During this upgrade you are asked if you want to update application support files as well. If a silent upgrade is performed by default, all installed support files are updated. 3. You can use the tacmd updateAgent command to install an agent update on a specified managed system. For reference information about this command and related commands, see the <i>IBM Tivoli Monitoring Command Reference</i>. 4. You must recycle (stop and restart) the portal server to use upgraded workspaces for your monitoring agents. 		

Linux and UNIX: Upgrading a portal server running as a non-root process

If the Tivoli Enterprise Portal Server is running as non-root on either Linux or AIX and you plan to upgrade it to IBM Tivoli Monitoring V6.3 using a non-root user ID, then you must perform the following procedure before upgrading.

On Linux, you need perform only “Step 1: Verify the DB2 Database for Linux, UNIX, and Windows authorizations”; on AIX, perform both “Step 1: Verify the DB2 Database for Linux, UNIX, and Windows authorizations” and “Step 2: Invoke the AIX slibclean command” on page 165.

Step 1: Verify the DB2 Database for Linux, UNIX, and Windows authorizations Before you begin

Both Linux and AIX sites must complete this step.

About this task

In the following example the portal server database is called **TEPS**, the DB2 for Linux, UNIX, and Windows userid is *itmuser*, and the DB2 administrative userid is *db2inst1*.

Procedure

1. Log in as the DB2 administrator:
`su - db2inst1`
2. Connect to the portal server database using the portal server's DB2 userid:
`db2 connect to TEPS user itmuser using TEPSpw`

where *TEPSpw* is the password for userid *itmuser*.

3. Get the authorizations for the portal server's DB2 userid:
`db2 get authorizations`

Example output:

Administrative Authorizations for Current User

Direct SYSADM authority	= NO
Direct SYSCTRL authority	= NO
Direct SYSMAINT authority	= NO
Direct DBADM authority	= NO
Direct CREATETAB authority	= NO
Direct BINDADD authority	= NO
Direct CONNECT authority	= NO
Direct CREATE_NOT_FENC authority	= NO
Direct IMPLICIT_SCHEMA authority	= NO
Direct LOAD authority	= NO
Direct QUIESCE_CONNECT authority	= NO
Direct CREATE_EXTERNAL_ROUTINE authority	= NO
Direct SYSMON authority	= NO

Indirect SYSADM authority	= NO
Indirect SYSCTRL authority	= NO
Indirect SYSMAINT authority	= NO
Indirect DBADM authority	= NO
Indirect CREATETAB authority	= YES
Indirect BINDADD authority	= YES
Indirect CONNECT authority	= YES
Indirect CREATE_NOT_FENC authority	= NO
Indirect IMPLICIT_SCHEMA authority	= YES

```

Indirect LOAD authority           = NO
Indirect QUIESCE_CONNECT authority = NO
Indirect CREATE_EXTERNAL_ROUTINE authority = NO
Indirect SYSMON authority         = NO

```

4. Ensure the settings listed below are set to YES, as shown.

```

Direct DBADM authority           = YES
Direct CREATETAB authority       = YES
Direct BINDADD authority         = YES
Direct CONNECT authority         = YES
Direct CREATE_NOT_FENC authority = YES
Direct IMPLICIT_SCHEMA authority = YES
Direct LOAD authority            = YES
Direct QUIESCE_CONNECT authority = YES
Direct CREATE_EXTERNAL_ROUTINE authority = YES

```

If they are all set to **YES**, you are finished. Otherwise continue with steps 5 through 7

5. Connect to the Tivoli Enterprise Portal Server database as the DB2 administrator:

```
db2 connect to TEPS user db2inst1 using db2pw
```

where *db2pw* is the password for userid *db2inst1*.

6. Grant the appropriate authorizations to the portal server's DB2 userid:

```
db2 GRANT DBADM,CREATETAB,BINDADD,CONNECT,CREATE_NOT_FENCED_ROUTINE,IMPLICIT_SCHEMA,LOAD,
CREATE_EXTERNAL_ROUTINE,QUIESCE_CONNECT ON DATABASE to user itmuser
```

7. Reconnect to the portal server database using the portal server's DB2 userid, and recheck its authorizations:

```
db2 connect to TEPS user itmuser using TEPSpw
Database Connection Information
```

```

Database server      = DB2/LINUX 8.2.8
SQL authorization ID = ITMUSER
Local database alias = TEPS

```

```
db2 get authorizations
Administrative Authorizations for Current User
```

```

Direct SYSADM authority           = NO
Direct SYSCTRL authority          = NO
Direct SYSMAINT authority         = NO
Direct DBADM authority            = YES
Direct CREATETAB authority        = YES
Direct BINDADD authority          = YES
Direct CONNECT authority          = YES
Direct CREATE_NOT_FENC authority   = YES
Direct IMPLICIT_SCHEMA authority   = YES
Direct LOAD authority             = YES
Direct QUIESCE_CONNECT authority   = YES
Direct CREATE_EXTERNAL_ROUTINE authority = YES
Direct SYSMON authority           = NO

Indirect SYSADM authority         = NO
Indirect SYSCTRL authority        = NO
Indirect SYSMAINT authority       = NO
Indirect DBADM authority          = NO
Indirect CREATETAB authority      = YES
Indirect BINDADD authority        = YES
Indirect CONNECT authority        = YES
Indirect CREATE_NOT_FENC authority = NO
Indirect IMPLICIT_SCHEMA authority = YES

```

```
Indirect LOAD authority           = NO
Indirect QUIESCE_CONNECT authority = NO
Indirect CREATE_EXTERNAL_ROUTINE authority = NO
Indirect SYSMON authority         = NO
```

You should now see the authorizations for the following items set to YES:

```
Direct DBADM authority           = YES
Direct CREATETAB authority       = YES
Direct BINDADD authority         = YES
Direct CONNECT authority         = YES
Direct CREATE_NOT_FENC authority = YES
Direct IMPLICIT_SCHEMA authority = YES
Direct LOAD authority            = YES
Direct QUIESCE_CONNECT authority = YES
Direct CREATE_EXTERNAL_ROUTINE authority = YES
```

Step 2: Invoke the AIX slibclean command

This step applies only to AIX sites, not to Linux sites.

As the *root* user, invoke the AIX `slibclean` command:

```
su -c "/usr/sbin/slibclean"
```

Required Java Runtime Environment

IBM Tivoli Monitoring V6.3 requires Java 7.

IBM Tivoli Monitoring agents on the same computer share the same Java environment, and will be automatically upgraded to the latest version when a V6.3 agent is installed.

Scenario: a rolling product upgrade

This section contains an example of a possible scenario for rolling out the upgrade of an installed IBM Tivoli Monitoring environment from one release to a later release.

- “Configuration”
- “Upgrading the Tivoli Monitoring environment” on page 166
- “Expected results” on page 167
- “Special instructions for reseeding a Hot Standby monitoring server” on page 168

Configuration

Your existing Tivoli Monitoring environment includes these five distributed servers.

- The primary (acting) hub Tivoli Enterprise Monitoring Server runs on server #1, an AIX 6.1, 64-bit node.
- The secondary (backup) hub monitoring server runs on server #2, a Solaris 10, 32-bit node.
- A remote Tivoli Enterprise Monitoring Server runs on server #3, an RHEL Linux 5, 32-bit node.
- A second remote monitoring server runs on server #4, a Windows 2008 node.
- The Tivoli Enterprise Portal Server running on server #5, an RHEL 6, 32-bit node, communicates with the acting (primary) hub monitoring server running on server #1. server #5 also runs the Warehouse Proxy Agent and the Summarization and Pruning Agent.

Each system runs an OS agent. All 14 nodes in this example configuration are connected to one remote monitoring server as its primary and the other as its secondary, to accommodate agent switching when upgrading the remote Tivoli Enterprise Monitoring Servers. Half of the agents point to each remote monitoring server as its primary monitoring server.

Upgrading the Tivoli Monitoring environment

Install a given level of IBM Tivoli Monitoring on the primary and secondary hub monitoring server, on the remote monitoring servers, and on the portal server. Install agents (simple agents, subnode agents, SYSPLEX agents), including the Warehouse Proxy Agent and the Summarization and Pruning Agent, at the same level of IBM Tivoli Monitoring.

Configure the Warehouse Proxy and Summarization and Pruning Agents to report directly into both hubs, configured as primary and secondary. You can optionally configure a second Warehouse Proxy Agent, to serve as a backup.

If you are upgrading from V6.1 or V6.2.x to V6.3, you must update both infrastructure components and OS agents in that \$ITM_HOME installation at the same time.

Other agents report to at least one of the remote monitoring servers.

1. Begin historical collection of agent attributes. Assume the historical collection is being done at the agents and not at the Tivoli Enterprise Monitoring Server; otherwise data collection may be interrupted when agents switch from the acting monitoring server to the backup.
2. Cause situations to fire, for example, by varying thresholds.
3. Ensure the acting (primary) and backup (secondary) monitoring servers are connected by examining their log/trace records and process state for these messages:
 - KQM0001 "FTO started at ..."
 - KQM0003 "FTO connected to ..."
 - KQM0009 "FTO promoted *primary* as the acting HUB"
 - KQM0009 "FTO promoted SITMON**secondary* (Mirror) as the acting HUB"
4. If you have integrated your IBM Tivoli Monitoring events with either Tivoli Enterprise Console or Netcool/OMNIBus:
 - a. First upgrade your event synchronization.
 - b. Then restart Tivoli Enterprise Console or OMNIBus, as appropriate.
 - c. Finally, restart the monitoring server to which Tivoli Enterprise Console or OMNIBus is connected. This must be the acting monitoring server.
5. Confirm that both the acting and backup hub monitoring servers are operational, so that during this upgrade process, when the acting hub is shut down, the backup will be able to assume the role of acting Tivoli Enterprise Monitoring Server.
6. Upgrade the primary hub Tivoli Enterprise Monitoring Server and all other Tivoli Management Services infrastructure components in that \$ITM_HOME installation (in other words, select ALL when asked what to upgrade).
 - If working on a Windows platform, apply support to upgraded agents as part of this upgrade step.
 - If working on a UNIX or Linux platform, apply support for the upgraded infrastructure agents on the primary hub monitoring server:

```
./itmcmd support -t primary_TEMS_name lz nt ux hd sy
```

Note: All Tivoli Management Services processes running within the same IBM Tivoli Monitoring environment are temporarily shut down when upgrading a monitoring server.

7. Upgrade the secondary hub Tivoli Enterprise Monitoring Server and all other Tivoli Management Services infrastructure components in that \$ITM_HOME installation (in other words, select ALL when asked what to upgrade).

Note: The acting (primary) and backup (secondary) monitoring servers should be installed at the same version level. The best choice is to have the monitoring servers at the same fix pack level as well.

- If working on a Windows platform, apply support to upgraded agents as part of this upgrade step.
 - If working on a UNIX or Linux platform, apply support for the upgraded infrastructure agents on the secondary hub monitoring server:
`./itmcmd support -t secondary_TEMS_name lz nt ux hd sy`
8. Restart both the primary (backup) and secondary (acting) Tivoli Enterprise Monitoring Servers.
 9. Upgrade the Tivoli Enterprise Portal Server.
 10. Confirm that all remote Tivoli Enterprise Monitoring Servers are operational, so that during this upgrade process, when the acting hub monitoring server is shut down, the backup will be able to assume the role of acting Tivoli Enterprise Monitoring Server.
 11. Upgrade each remote monitoring server, one at a time.
 - If working on a Windows platform, apply support to upgraded agents when upgrading the monitoring server.
 - If working on a UNIX or Linux platform, immediately after a remote monitoring server is upgraded, apply support for the upgraded infrastructure agents on the same node. Then restart the remote monitoring server to make the upgraded support effective.
 12. Upgrade the remaining infrastructure agents.
 13. At a time when no history records will be uploaded to it, upgrade the Warehouse Proxy Agent.

Expected results

This section contains the expected results of the upgrade scenario.

1. When the primary hub monitoring server is down while completing step 6:
 - a. Configure the portal server to point to the newly acting hub monitoring server, and verify that expected agents and situations show up on the Tivoli Enterprise Portal client.
 - b. Verify that the secondary monitoring server has taken over as acting hub, in other words, that these messages appear in its log:
 - KQM0004 "FTO detected lost parent connection"
 - KQM0009 "FTO promoted *secondary* as the acting HUB"
 - c. Verify that the event handler your site is using (either Tivoli Enterprise Console or Netcool/OMNIbus) is still being updated, now by the secondary hub, by causing an event state to change and validating the state of all events on the remote console.
 - d. Verify that attribute collection was uninterrupted by examining the history data, by varying attribute values at the agent, and by observing the change in the portal client.
 - e. Ensure that the primary hub Tivoli Enterprise Monitoring Server is not restarted by the upgrade process until all remote monitoring servers and directly connected agents have successfully failed over to the secondary hub.
2. After completing step 6, verify that the primary hub Tivoli Enterprise Monitoring Server has taken on the role of standby monitoring server: ensure its log contains these messages:
 - KQM0001 "FPO started at ..."
 - KQM0003 "FTO connected to ..."
 - KQM0009 "FTO promoted *secondary* as the acting HUB"
 - KQM0009 "FTO promoted *primary*(Mirror) as the acting HUB"

The log for the secondary hub monitoring server should contain these messages:

- KQM0005 "FTO has recovered parent connection ..."
- KQM0009 "FTO promoted *secondary* as the acting HUB"

Use the sitpad and taudit tools for assistance. These tools can be downloaded from the IBM Tivoli Integrated Service Management Library Web site at IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>).

3. When the secondary hub Tivoli Enterprise Monitoring Server is stopped while completing step 7:
 - a. Configure the Tivoli Enterprise Portal Server to point to the acting hub, the primary.

- b. Verify that the primary monitoring server has taken over as the acting hub; its log should contain these messages:
 - KQM0004 "FTO detected lost parent connection"
 - KQM0009 "FTO promoted *primary* as the acting HUB"
- c. Verify that either Tivoli Enterprise Console or Netcool/OMNIBus is still being updated with event information.
- d. Verify that attribute history is still being collected as in substep 1.d.
- e. Ensure that the secondary hub Tivoli Enterprise Monitoring Server is not restarted by the upgrade process until all remote monitoring servers and directly connected agents have successfully failed over to the primary hub.

4. After completing step 7, verify that the secondary hub Tivoli Enterprise Monitoring Server is functional by examining the log/trace records of both hub monitoring servers for these messages.

In the primary hub's log:

- KQM0005 "FTO has recovered parent connection ..."
- KQM0009 "FTO promoted *primary* as the acting HUB"

In the secondary hub's log:

- KQM0001 "FPO started at ..."
- KQM0003 "FTO connected to ..."
- KQM0009 "FTO promoted *primary* as the acting HUB"
- KQM0009 "FTO promoted *secondary*(Mirror) as the acting HUB"

Use the sitpad and taudit tools for assistance.

5. When the Tivoli Enterprise Monitoring Server is stopped while completing step 11 for a given remote monitoring server, verify that either Tivoli Enterprise Console or Netcool/OMNIBus is still being updated, now by the remote monitoring server configured for each agent, and that attribute collection was uninterrupted, as in substep 1.d. Examine the log/trace records of each remote Tivoli Enterprise Monitoring Server to verify agent reporting; use the sitpad and taudit tools for assistance.
6. After completing step 11 for a given remote Tivoli Enterprise Monitoring Server, verify that the monitoring server is operational by examining its log/trace records.

Note: Agents do not automatically switch back to their primary monitoring server unless the secondary server fails or the agents are manually restarted.

7. Ensure the historical collection of attributes is unbroken, in both short-term and long-term history.
8. Ensure situations are reported to Tivoli Enterprise Console and to Netcool/OMNIBus without interruption. Verify that always-true situations are still true after agent switching.
9. Verify that the Tivoli Enterprise Portal client for the acting hub Tivoli Enterprise Monitoring Server shows attribute and situation data at all times.

Note:

1. Pure events are lost when failover occurs from the primary hub to the secondary and back, as events are not replicated across the two hubs.
2. All situations are restarted during failover, so a situation that was already raised at the primary hub will be raised once again if the situation is still true when failover occurs.
3. A master reset is issued to Tivoli Enterprise Console when the failover occurs, with this event: "Hotstandby TEMS switched. New Primary TEMS *hostname_of_newly_active_TEMS*." Events raised after the failover will be resent.
4. When a failing monitoring server is restarted, Tivoli Enterprise Console receives this event: "TEMS *hostname_of_restarted_TEMS* restarted."

Special instructions for reseeding a Hot Standby monitoring server

Follow the steps in this section if your site has implemented a Hot Standby monitoring server with agent switching and you want to ensure a monitoring server remains active at all times, even while upgrading

its application support (remember that upgrading any IBM Tivoli Monitoring process shuts down all local Tivoli Management Services components while the upgrade takes place).

Before you begin

This procedure assumes your site needs to keep a Hot Standby Tivoli Enterprise Monitoring Server running at all times. If that is not necessary, it is simpler to shut down the standby monitoring server until after seeding takes place, then start it back up.

Procedure

1. Upgrade one of the two hubs, either the active one or the backup. During this process, the other hub monitors your IBM Tivoli Monitoring environment.
2. Upgrade the other hub, during which time the newly upgraded hub monitors your environment.
3. With both upgraded hubs running, one in active mode and the other in backup mode, reseed the active hub without restarting it.
4. Reseed the backup hub.
5. Restart the hub you want to be active, during which time the other hub monitors your IBM Tivoli Monitoring environment.
6. Restart the hub you want to be in backup mode, which puts the other hub in active mode and this hub in backup mode.

Chapter 9. Installing IBM Tivoli Monitoring on one computer

This chapter describes how to deploy a small IBM Tivoli Monitoring environment on a single Windows computer. Installation on one computer might be useful for a test environment, a teaching environment, or for monitoring a small server environment.

The single-computer installation supports between 100 and 200 monitoring agents and can perform minimal historical data collection. See “Sizing your Tivoli Monitoring hardware” on page 34 for further discussion of the capabilities of a small monitoring environment.

The single-computer installation includes the following components:

- The hub Tivoli Enterprise Monitoring Server (hub monitoring server)
- The Tivoli Enterprise Monitoring Automation Server
- The Tivoli Enterprise Portal Server (portal server)
- The Tivoli Enterprise Portal Client (portal client)
- The Tivoli Enterprise Portal Server database (portal server database)
- The Tivoli Data Warehouse database (warehouse database)
- The Warehouse Proxy Agent
- The Summarization and Pruning Agent
- The Tivoli Performance Analyzer

The single-computer installation does not include event forwarding to either IBM Tivoli Enterprise Console or IBM Tivoli Netcool/OMNIbus. There is no user authentication for the monitoring server or portal server.

If you are using Jazz for Service Management components, review the hardware and software requirements in the Jazz for Service Management Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/psc_ic-homepage.html) to determine if your computer has sufficient hardware to support installing IBM Tivoli Monitoring and Jazz for Service Management together. You can install them on a single computer that meets the hardware requirements for a small environment. However, it is recommended that you install Jazz for Service Management on a separate computer from your single-computer installation of IBM Tivoli Monitoring for a production environment.

Prerequisites for the single-computer installation

See the software prerequisites for the single-computer installation described in this chapter.

- A supported Windows operating system. A 64 bit Windows operating system is recommended.
- If your site does not plan to use or cannot use (due to capacity limitations) the embedded Derby database, you will also need one of the following RDBMS servers for the portal server:
 - IBM DB2 Database for Linux, UNIX, and Windows
 - Microsoft SQL Server

Note: If you plan to use either of the listed RDBMSes with the portal server, it must already have been installed and be running when you begin this installation. If you plan to use the Derby database, it is not necessary that you preinstall it; this embedded RDBMS is installed, if necessary, with IBM Tivoli Monitoring.

- If you want IBM Tivoli Performance Analyzer to integrate with IBM SPSS Statistics Server, it must already be installed and running when you begin this installation.

For the Tivoli Data Warehouse, you can also use Microsoft SQL Server. A 64-bit version of IBM DB2 for Linux, UNIX, and Windows Enterprise Server Edition 10.1 is included with IBM Tivoli Monitoring. This version should be used for new installations to minimize the need to upgrade DB2 for Linux, UNIX, and Windows in the future.

While it is possible to install Tivoli Monitoring on a single Linux or UNIX computer, additional manual configuration is required. When you install Tivoli Monitoring on a Windows computer with IBM DB2 for Linux, UNIX, and Windows or Microsoft SQL Server, much of the configuration required to set up the Tivoli Data Warehouse is automated. (An extra configuration step is required after installation if you use Microsoft SQL Server for the Tivoli Data Warehouse database.)

In a single-computer installation, the RDBMS server is used for both of the databases required by Tivoli Monitoring:

- The *Tivoli Enterprise Portal Server database* (or *portal server database*) stores user data and information required for graphical presentation on the user interface. The portal server database is created automatically during configuration of the portal server. It is always located on the same computer as the portal server.
- The *Tivoli Data Warehouse database* (also called the *warehouse database* or *data warehouse*) stores historical data for presentation in historical data views. In a single-computer installation, the warehouse database is created on the same relational database management server (RDBMS) used for the portal server database. In larger environments, it is best to create the warehouse database on a different computer from the portal server.

Installation procedure

Complete the steps in this section to install IBM Tivoli monitoring on one Windows computer.

Procedure

1. Launch the installation wizard by double-clicking the setup.exe file in the WINDOWS subdirectory of the installation media.
2. Click **Next** on the Welcome window.
3. Click **Accept** to accept the license agreement.
4. Choose the directory where you want to install the product. The default directory is C:\IBM\ITM. Click **Next**.

Note: If you specify an incorrect directory name, you will receive the following error:

The IBM Tivoli Monitoring installation directory cannot exceed 80 characters or contain non-ASCII, special or double-byte characters.

The directory name can contain only these characters:

"abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ _\.:0123456789()~-. /".

5. Click **Next** to accept the default encryption key and then click **OK** on the popup window to confirm the encryption key.
6. On the Select Features window, select the check boxes for the components you want to install. Select all components for a complete installation on one computer.

For additional information about these components, press the **Help** button.

Note: The agent support comprises additional support files for the Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, and the Tivoli Enterprise Portal desktop client; these are automatically installed whenever you install the agent support. Also, the Eclipse Help Server is automatically installed whenever you install a Tivoli Enterprise Portal Server.

7. Click **Next**.

The Agent Deployment window is displayed. This window lists monitoring agents that you can deploy to remote computers. For this installation, do not select any agents for remote deployment.

8. Click **Next**.
9. If the TEPS Desktop and Browser Signon ID and Password window is displayed, enter and confirm the password to be used for logging on to the Tivoli Enterprise Portal. The default logon user ID, **sysadmin**, cannot be changed on this window. This password is required only when Security: Validate User has been enabled on the hub monitoring server. This window is not displayed if the **sysadmin** user ID has already been defined in the operating system.
10. Review the installation summary details. The summary identifies the components you are installing. Click **Next** to begin the installation. The status bar indicates the progress of your installation.
After the components are installed, a configuration window is displayed.
11. Click **Next** to start configuring all selected components. (Leave all check boxes selected for this installation.)
12. Configure communications for the Tivoli Enterprise Portal Server:
 - a. Click **Next** to confirm that you are installing the portal server on this computer. (The host name of this computer is displayed by default.)
 - b. If no relational database manager can be found in this computer, the embedded Derby RDBMS is used by default. However, if at least one RDBMS product is installed on this computer, a window is displayed for you to choose the RDBMS product you want to use. Choose either **Embedded TEPS database** (that is, Derby), **IBM DB2 Universal Database™ Server**, or **Microsoft SQL Server**, and click **Next**.
 - c. A window is displayed for you to configure the connection between the portal server and the *portal server database* (TEPS database).
The installation program uses the information on this window to automatically perform the following tasks:
 - Create the portal server database.
 - Create a database user for the portal server to use to access the database.
 - Configure the ODBC connection between the portal server and the database.

TEPS Data Source Config Parameters - DB2

Data source name:
TEPS2

Please enter your Database Administrator ID (up to 8 characters) and password below:

Admin User ID:
db2admin

Admin Password:
xxx

Please accept the default Database User ID or enter your own TEPS Database User ID (up to 8 characters) and password:

Database User ID:
TEPS

Database Password:
xxx

Reenter Password:
xxx

Please select database instance name from list below. If you leave it empty default instance will be selected.

OK Cancel Help

Figure 17. Configuration window for the portal server database using DB2 for Linux, UNIX, and Windows

Figure 17 shows the configuration window for a portal server database using DB2 for Linux, UNIX, and Windows. The configuration window for a Derby or Microsoft SQL Server database is similar. The fields on the configuration window are described in the following table:

Table 29. Configuration information for the portal server database

Field	DB2 for Linux, UNIX, and Windows default	MS SQL default	Description
Admin User ID	db2admin	sa	The database administrator ID.
Admin Password	(no default)	(no default)	The password for the database administrator ID.
Database User ID	TEPS	TEPS	The login name of the database user that the portal server will use to access the database.
Database Password	(no default)	(no default)	The password for the database login user. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Reenter password	(no default)	(no default)	Confirm the password by entering it again.
Select database instance name	(no default)	(no default)	Pull down the drop-down list, and select the appropriate database instance.

- d. Optionally change the default values for the administrator ID and database user ID. Enter passwords for the administrator and database user. Click **OK**.

- e. Click **OK** on the message that tells you that the portal server configuration was successful.
- f. Click **Next** to accept the default Tivoli Data Warehouse user ID and password. (You can change these values in a later step.)
- g. In the TEP Server Configuration window, click **OK** to accept the default communications protocol, **IP.PIPE**. IP.PIPE is the protocol that the portal server will use to communicate with the Tivoli Enterprise Monitoring Server.
A second TEP Server Configuration window is displayed. The **IP.PIPE** area of this window displays the host name of this computer and the default port number, 1918.
- h. Do not select **Validate user with LDAP** since user authentication is not enabled during a single computer installation. It can be configured at a later time.
- i. Select the **Enable the dashboard data provider** check box if Dashboard Application Services Hub will be used with monitoring dashboard applications such as IBM Infrastructure Management Dashboards for Servers or will be used to create custom dashboard pages for monitoring agents. See the Preparing your dashboard environment topic in the *IBM Tivoli Monitoring Administrator's Guide* for details on additional configuration tasks to perform after installation.
- j. Do not specify a value for the domain override or enable authorization policies. Authorization policies are not enabled during a single computer installation. They can be configured at a later time.
- k. Click **OK** to accept the default host name and port number.
- l. Click **Yes** on the message asking if you want to reconfigure the warehouse connection information for the portal server.
- m. Select **DB2** or **SQL Server** from the list of RDBMS platforms and click **OK**.
A window is displayed for you to configure the connection between the portal server and the *Tivoli Data Warehouse database*. (The Tivoli Data Warehouse database is referred to in the title of this window as the data source for the Warehouse Proxy. The Warehouse Proxy Agent sends information collected from monitoring agents to the Tivoli Data Warehouse.)
The installation program uses the information on this window to automatically perform the following tasks:
 - Create the Tivoli Data Warehouse database.
 - Create a database user (called the *warehouse user*) for the portal server, Warehouse Proxy Agent, and Summarization and Pruning Agent to use to access the warehouse database.
 - Configure the ODBC connection between the portal server and the warehouse database.

Figure 18. Configuration window for the Tivoli Data Warehouse database using DB2 for Linux, UNIX, and Windows

Figure 19. Configuration window for the Tivoli Data Warehouse database using Microsoft SQL Server

The fields on this window are described in the following table:

Table 30. Configuration information for the Tivoli Data Warehouse database

Field	DB2 for Linux, UNIX, and Windows default	MS SQL default	Description
Data Source Name	ITM Warehouse	ITM Warehouse	The name of the data source. ^{1 2}
Database Name	WAREHOUS	WAREHOUS	The name of the database.
Admin User ID	db2admin	sa	The database administrator ID.
Admin Password	(no default)	(no default)	The password for the database administrator ID.
Database User ID	ITMUser	ITMUser	The login name of the database user that the portal server, Warehouse Proxy Agent, and Summarization and Pruning Agent will use to access the Tivoli Data Warehouse database.
Database Password	itmpswd1	itmpswd1	The password for the database login user. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Reenter password	itmpswd1	itmpswd1	Confirm the password by entering it again.

Note:

- 1) Warehouse Proxy Agent does not create this DSN - it must already exist. If you are installing Performance Analyzer on the same machine where the Tivoli Enterprise Portal Server is installed, you can use the existing data source created by Tivoli Monitoring. Otherwise, you must create a new System DSN manually, prior to re-configuring Performance Analyzer.

- 2) On 64-bit versions of Windows, data sources created by the default **ODBC Data Source Administrator** applet available from the **Control Panel** are not available if you are installing the 32-bit version of the portal server. In this scenario, you must use the 32-bit version of the **ODBC Data Source Administrator** applet from <WINDOWS>\SysWOW64\odbcad32.exe. If you are installing the 64-bit version of the portal server then you can use data source created by the ODBC Data Source Administrator applet.
 - n. Optionally change the default values on this window. Enter the password of the database administrator. Click **OK**.
 - o. Click **OK** on the message that tells you that the portal server configuration was successful.
 13. If your monitoring server is not local, you must configure communications for the hub Tivoli Enterprise Monitoring Server:
 - a. Remove the check mark from **Security: Validate User**. User authentication is not enabled during a single computer installation. It can be configured at a later time.
 - b. Enabling **TEMS Security Compatibility Mode** performs the following functions:
 - 1) It allows for backwards compatibility with previous versions of IBM Tivoli Monitoring components shipped before V6.3 when securing command execution (Situation Take Action, Execute Command, Take Action, Workflow Policy Take Action). Accept the default if you have a previous version of a portal server, hub monitoring server, or OMEGAMON product connecting to this Tivoli Enterprise Monitoring Server or if the system clocks are not synchronized across the systems with monitoring servers, the portal server, and monitoring agents.
 - 2) It permits commands to still execute even if cryptographic facilities are not available. This gives you time to ensure that, for example, your z/OS Tivoli Enterprise Monitoring Server has the Integrated Cryptographic Services Facility (ICSF) available.
- Note:** If you disable this option and you have a mixture of pre-V6.3 and V6.3 server components, if clocks are not synchronized, or if you don't have cryptographic services available, then command executions might fail.
- c. Click **OK** to accept the other defaults on the Tivoli Enterprise Monitoring Server Configuration window:
 - The type of server you are configuring is **Hub**.
 - The default name of the hub monitoring server is `HUB_host_name`, for example, `HUB_ITMSERV16`.
 - The default communications protocol is **IP.PIPE**. IP.PIPE is the protocol that the monitoring server will use to communicate with the portal server.
If your site uses the IP.PIPE or IP.SPIPE protocols for communications, be aware of the port usage rules for Tivoli Monitoring components. See "Controlling port number assignments" on page 421 for more information.
 - Do not enable user authentication until you have completed all installation and configuration and verified that the product is working. If you want to enable user authentication, see the *IBM Tivoli Monitoring Administrator's Guide*.

The Hub TEMS Configuration window is displayed. The **IP.PIPE** area of this window displays the host name of this computer and the default port number, 1918.
 - d. Click **OK** to accept the default host name and port number.
 14. You are asked if you want to configure the Tivoli Enterprise Monitoring Automation Server. If you select **Yes** the **Configure Tivoli Enterprise Monitoring Automation Server** window is displayed. Here you complete your Registry Services configuration details and Tivoli Enterprise Monitoring Automation Server connection details. When you select **Configure Registry Services connection**, you must specify the correct Protocol, Hostname, and Port number for the WebSphere server where Registry Services is installed.

The **User ID** and **Password** fields must be completed with the User ID and Password values for your Registry Services provider registry. The following criteria apply to the User ID and Password fields:

- Providing both a user ID and a password is accepted.
- Leaving both the user ID and a password fields blank is accepted.
- Providing one parameter without the other is not accepted.

In the **Automation Server connection information** section, provide the HTTP protocol and port for clients to connect to the automation server.

- If your monitoring server is not local, you must configure the default communication between any IBM Tivoli Monitoring component and the hub Tivoli Enterprise Monitoring Server:
 - On the configuration window, click **OK** to accept the default communications protocol, **IP.PIPE**. A second configuration window is displayed. The **IP.PIPE** area of this window displays the host name of this computer and the default port number, 1918.
 - Click **OK** to accept the default host name and port number.
- The **Configure Tivoli Performance Analyzer** window is displayed. Configure the agent for Data Warehouse. Some predefined values will appear depending on whether you choose ODBC or JDBC from the **Agent Database Connection Type** dropdown list.

Figure 20. Configure Tivoli Performance Analyzer window - ODBC

If you choose ODBC from the *Agent Database Connection Type* dropdown list, you must specify the following:

- Set your **Database Type** to DB2, MSSQL or ORACLE
- Specify the Data Source Name - **Agent ODBC DSN** (ITM Warehouse by default)

Note:

- Tivoli Performance Analyzer does not create this DSN - it must already exist. If you are installing Performance Analyzer on the same machine where the Tivoli Enterprise Portal Server

- is installed, you can use the existing data source created by Tivoli Monitoring. Otherwise, you must create a new System DSN manually, prior to re-configuring Performance Analyzer.
- On 64-bit versions of Windows, data sources created by the default **ODBC Data Source Administrator** applet available from the **Control Panel** are not available for 32-bit applications. Therefore you must use the 32-bit version of the **ODBC Data Source Administrator** applet from <WINDOWS>\SysWOW64\odbcad32.exe.
- c. Type the **Username** and **Password**. The entries in these fields are used to connect to the Tivoli Data Warehouse and are the same credentials as those used by the Tivoli Enterprise Portal Server, the Warehouse Proxy Agent and the Summarization and Pruning Agent to communicate with Tivoli Data Warehouse.

Figure 21. Configure Tivoli Performance Analyzer window - JDBC

If you choose JDBC from the Agent Database Connection Type dropdown list, you need to specify the following:

- Set your Database Type to DB2, MSSQL or ORACLE.
- Provide the **Hostname**. The Hostname is the host where the Tivoli Data Warehouse is installed.
- Provide the **Port** number.
- Provide the **Database Name**.
- Type the **Username** and **Password**.
- Specify the JDBC Driver. The default driver name for DB2 is com.ibm.db2.jcc.DB2Driver.
- Specify the JDBC Driver Path, which should be provided as a list of JAR files with the full path separated by ";".

Note: You can use the **Browse** button in order to specify the path. If you use the **Browse** button a file list is added at the end of the **JDBC Driver Path** text field separated from the existing content by a path separator. The data is used to initialize the database, so it must be provided no

matter what Agent Database Connection Type was selected. The driver path should be provided as a list of JAR files with the full path separated by ";".

17. Regardless of whether you choose ODBC or JDBC as your **Agent Database Connection Type** you can use the **Test connection** button to check whether the connection can be initiated. Pressing OK button launches more tests. Finishing the installation with incorrect connection settings is impossible unless you use the **Bypass connection tests** option in the **Advanced configuration** window.
18. Click **Next** to proceed to the **Advanced Configuration** window.

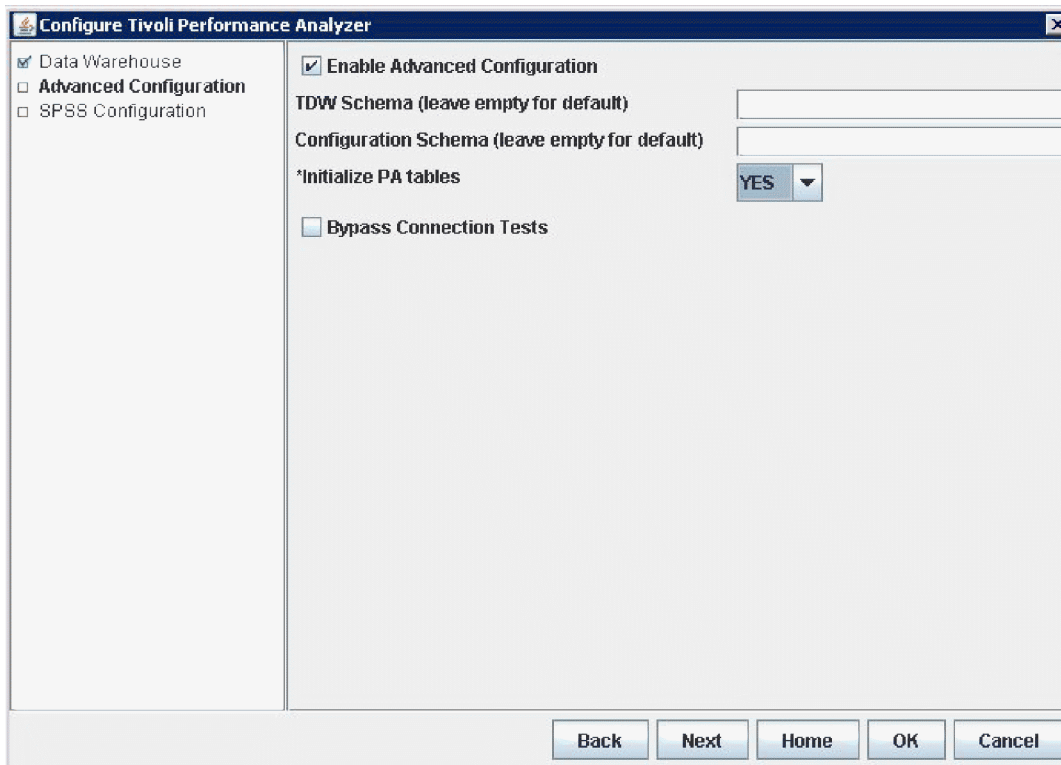


Figure 22. Configure Tivoli Performance Analyzer window - Advanced configuration

Note: Do not change any of the values displayed in this window unless you are an advanced user. Unless you select **Enable advanced configuration**, all options are greyed out.

- You can enable Advanced Configuration in order to specify **TDW Schema** and **Configuration Schema**.
- You can also choose to have the agent to Initialize PA tables.

Note: Setting **Initialize PA tables** to **YES** removes and recreates all previously created tables deleting all user tasks and reverting each OS task to its default. If you are upgrading from a previous installation of Tivoli Monitoring, set **Initialize PA tables** to **NO**.

- Use the **Bypass connection tests** option to finish the installation without running connection tests.

19. You can configure Performance Analyzer to use the IBM SPSS Statistics Server to support non-linear trending and forecasting of capacity and performance metrics. To enable SPSS, select **Enable SPSS configuration**.

Note: The IBM SPSS Statistics Server is a prerequisite for using the non-linear trending feature in Performance Analyzer. SPSS Statistics Server documentation is available at <http://publib.boulder.ibm.com/infocenter/spssstat/v20r0m0/index.jsp>, or from the Passport Advantage website with part number CI211EN, <http://www-01.ibm.com/software/howtobuy/>

passportadvantage/.

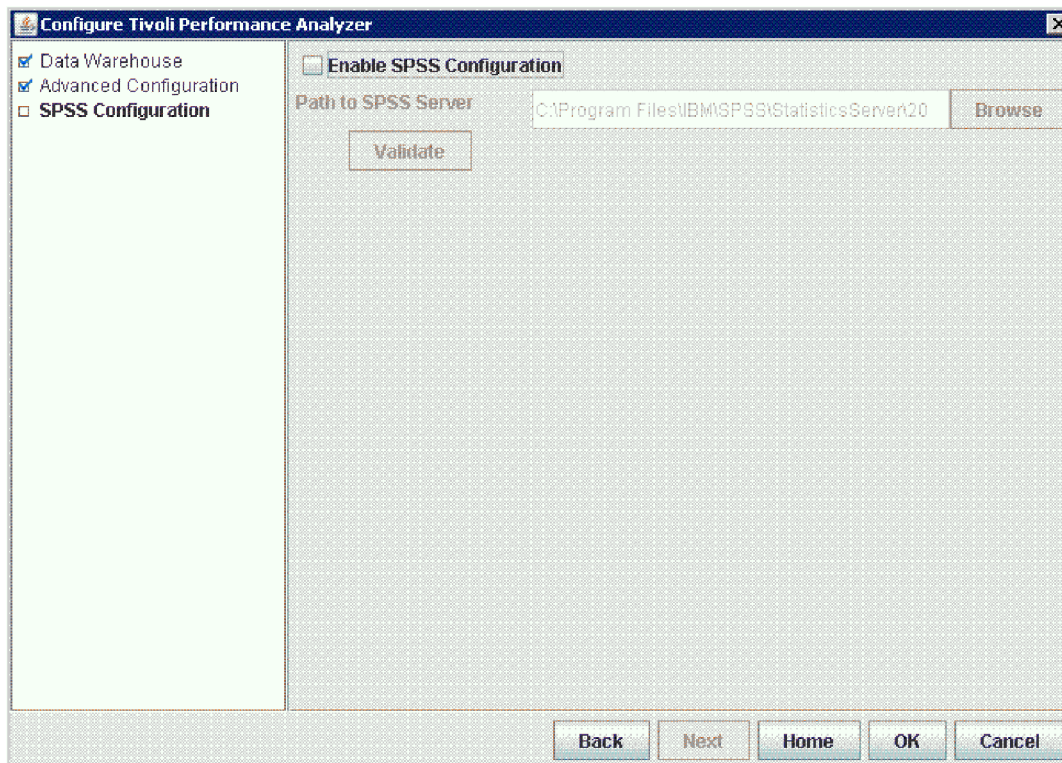


Figure 23. Configure Tivoli Performance Analyzer SPSS Configuration

20. Browse to the location of your SPSS Statistics Server. Click **Validate** to confirm the path you entered is correct.

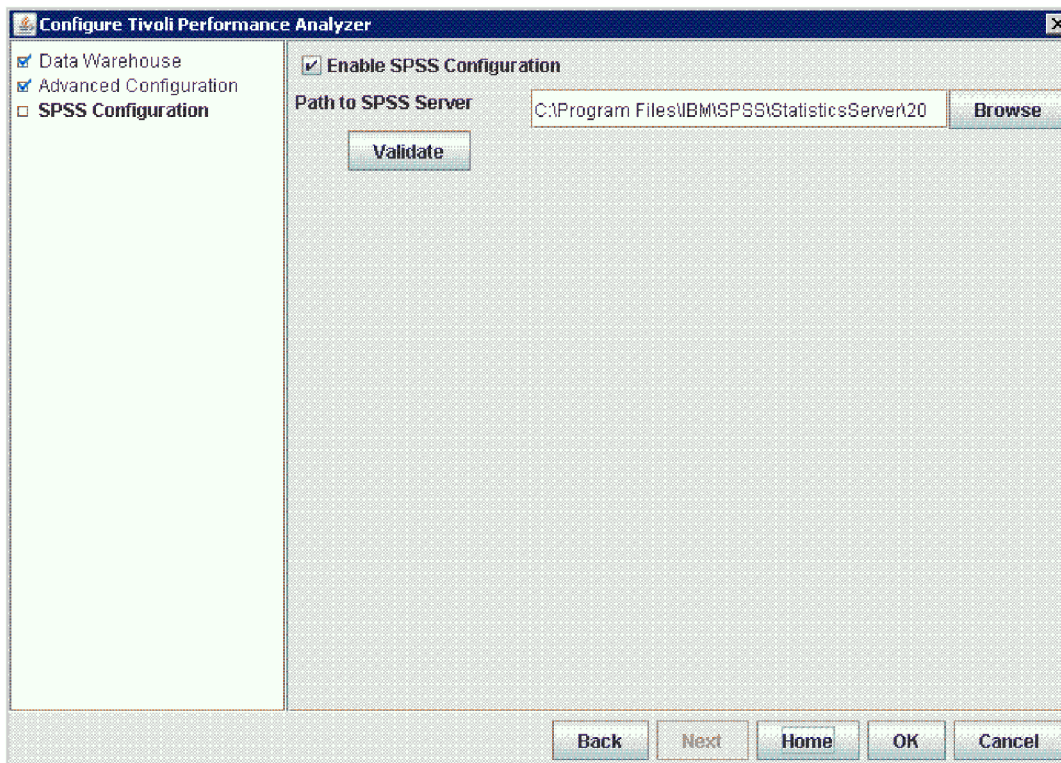


Figure 24. Configure Tivoli Performance Analyzer Enable SPSS Configuration

21. Click OK.

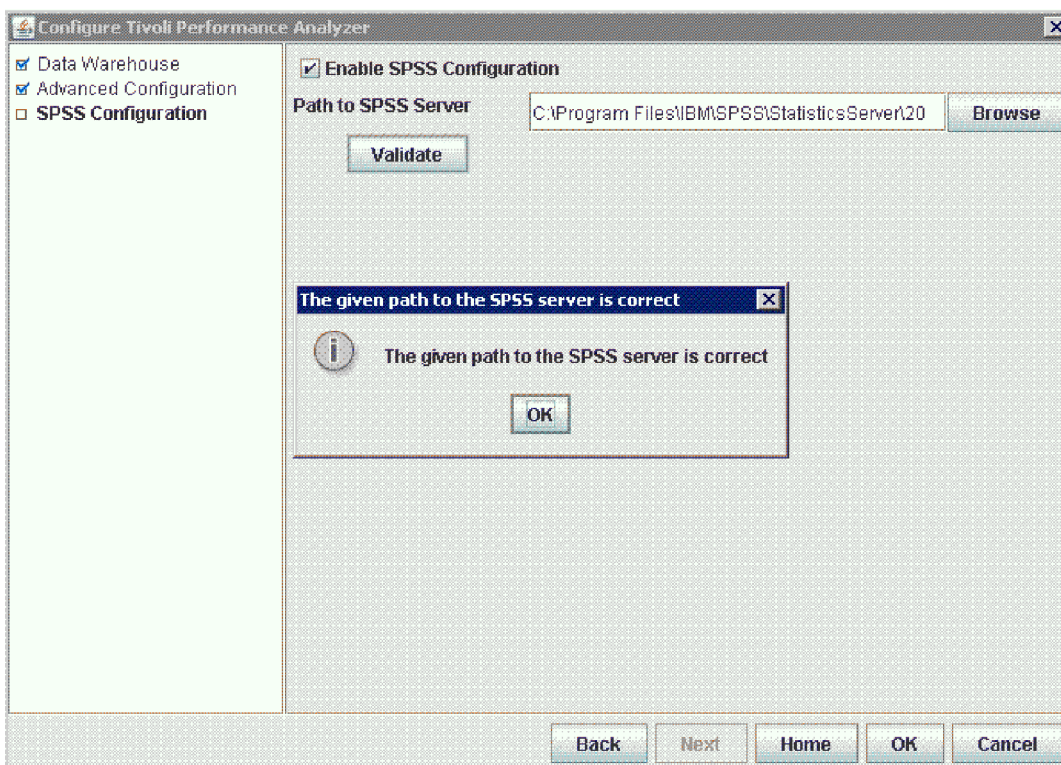


Figure 25. Configure Tivoli Performance Analyzer

22. Click **OK**. The system configures and starts the agent.
23. Click **Yes** on the message asking if you want to configure the Summarization and Pruning Agent.
A multi-tabbed configuration window is displayed with the **Sources** tab at the front.
Figure 26 shows the configuration window for a Summarization and Pruning Agent on Windows (values displayed are for a DB2 for Linux, UNIX, and Windows warehouse database). The configuration window for a Summarization and Pruning Agent on Linux or UNIX is similar.

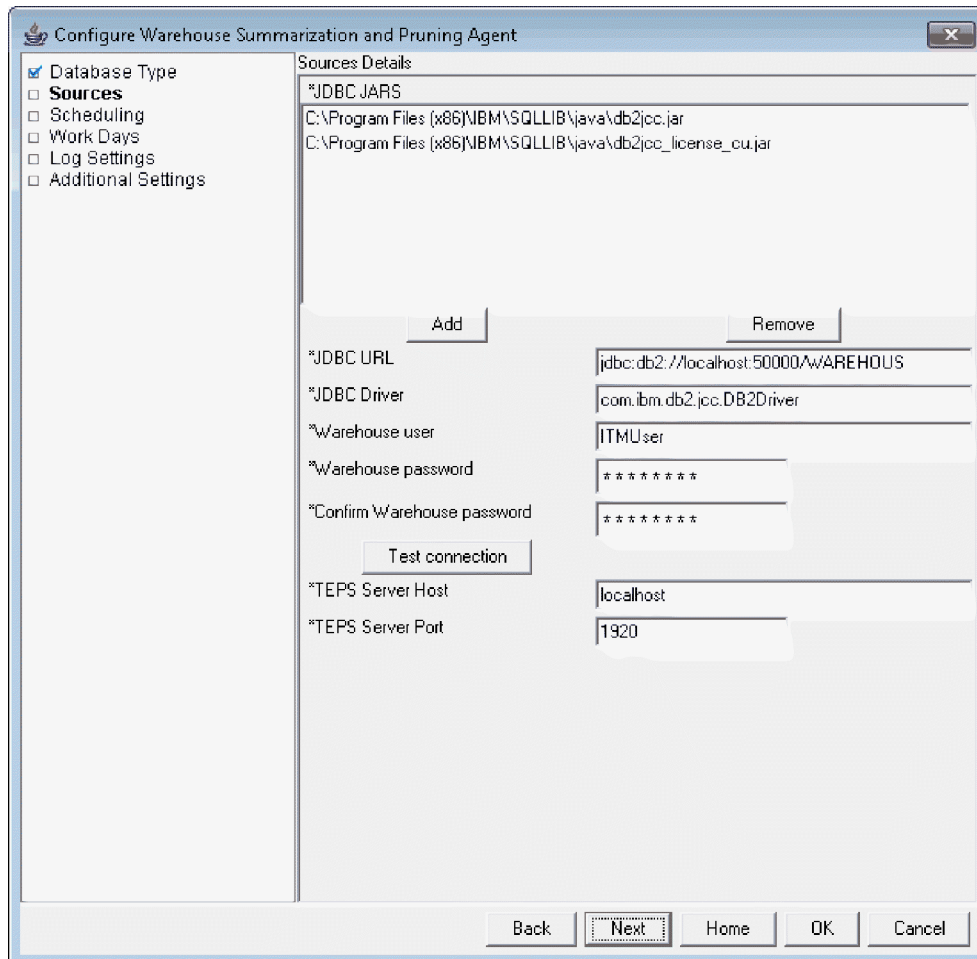


Figure 26. Sources pane of Configure Summarization and Pruning Agent window

24. Add the names and directory locations of the JDBC driver JAR files to the **JDBC Drivers** list box:
 - a. Click **Add** to display the file browser window. Navigate to the location of the driver files on this computer and select the Type 4 driver files for your database platform. See for the names and default locations of the driver files to add.
 - b. Click **OK** to close the browser window and add the JDBC driver files to the list.
 If you need to delete an entry from the list, select the entry and click **Remove**.
25. The default values for the database platform you selected in the Database Type pane are displayed in the other text fields on the **Sources** pane. Change the default value displayed in the **JDBC URL** field if it is not correct. The following table lists the default Tivoli Data Warehouse URLs for the different database platforms:

Table 31. Tivoli Data Warehouse URLs

Database platform	Warehouse URL
IBM DB2 for Linux, UNIX, and Windows	jdbc:db2://localhost:60000/WAREHOUS
Oracle	jdbc:oracle:thin:@localhost:1521:WAREHOUS
Microsoft SQL Server	jdbc:sqlserver://localhost:1433;databaseName=WAREHOUS

- If the Tivoli Data Warehouse is installed on a remote computer, specify the host name of the remote computer instead of localhost.
- Change the port number if it is different.
- If the name of the Tivoli Data Warehouse database is not WAREHOUS, replace WAREHOUS with the actual name.

26. Verify the JDBC driver name.

The following table lists the JDBC Type 4 driver names for each database platform:

Table 32. JDBC driver names

Database platform	JDBC driver name
IBM DB2 for Linux, UNIX, and Windows	com.ibm.db2.jcc.DB2Driver
Oracle	oracle.jdbc.driver.OracleDriver
Microsoft SQL Server	com.microsoft.sqlserver.jdbc.SQLServerDriver

27. If necessary, change the entries in the **Warehouse user** and **Warehouse password** fields to match the user name and password that were created for the Tivoli Data Warehouse. The default user name is itmuser and the default password is itmpswd1.
28. In the **TEPS Server Host** and **TEPS Server Port** fields, enter the host name of the computer where the Tivoli Enterprise Portal Server is installed and the port number that it uses to communicate with the Summarization and Pruning Agent.

Note: The default Tivoli Enterprise Portal Server interface port of 15001 is also used after the Summarization and Pruning Agent's initial connection to the portal server over port 1920. Any firewalls between the two need to allow communications on either 15001 or whichever port is defined for any new Tivoli Enterprise Portal Server interface used per the instructions in "Defining a Tivoli Enterprise Portal Server interface on Windows" on page 876.

29. Click **Test connection** to ensure you can communicate with the Tivoli Data Warehouse database.
30. Select the **Scheduling** check box to specify when you want summarization and pruning to take place. You can schedule it to run on a fixed schedule or on a flexible schedule:

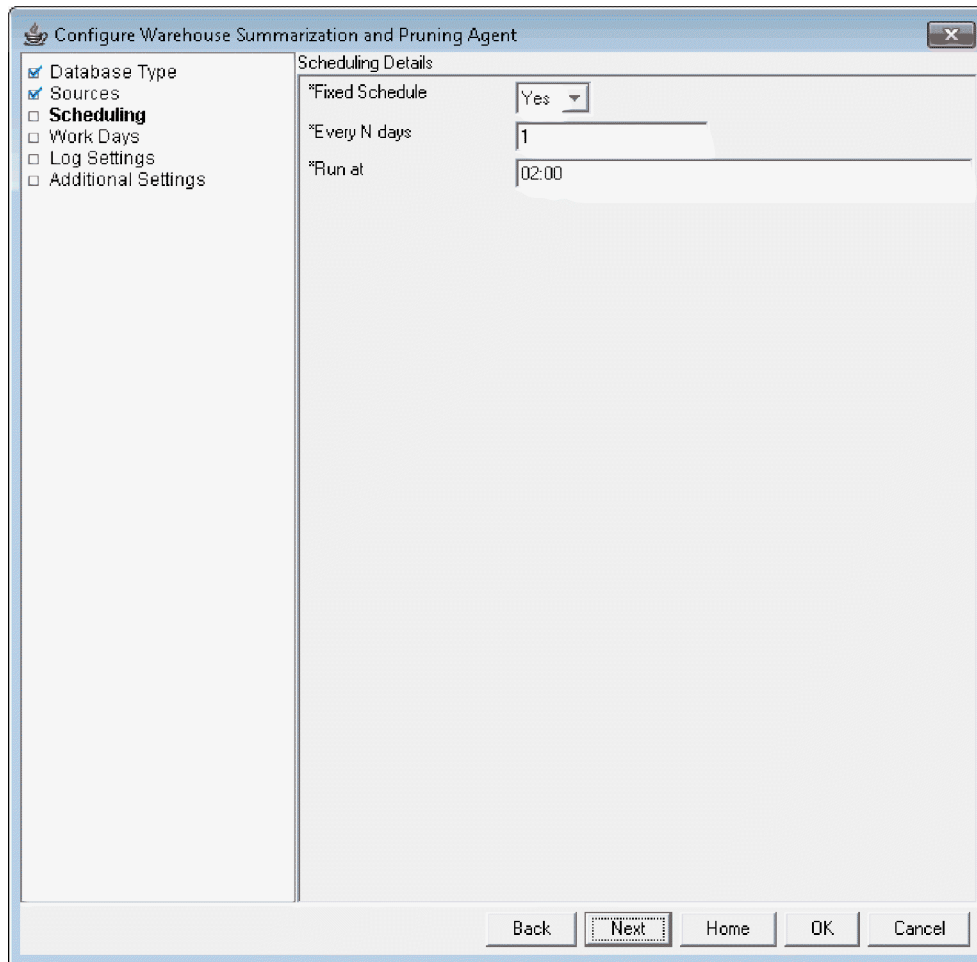


Figure 27. Scheduling pane of Configure Summarization and Pruning Agent window

Note: If you select Fixed, the Summarization and Pruning Agent does not immediately perform any summarization or pruning when it *starts*. It performs summarization and pruning when it *runs*. It runs according to the schedule you specify on the **Scheduling** pane. If you select Flexible, the Summarization and Pruning Agent runs once immediately after it is started and then at the interval you specified except during any blackout times.

31. Specify shift and vacation settings in the **Work Days** pane:




Figure 28. Work Days pane of Summarization and Pruning Agent configuration window

When you enable and configure shifts, IBM Tivoli Monitoring produces three separate summarization reports:

- Summarization for peak shift hours
- Summarization for off-peak shift hours
- Summarization for all hours (peak and off-peak)

Similarly, when you enable and configure vacations, IBM Tivoli Monitoring produces three separate summarization reports:

- Summarization for vacation days
- Summarization for nonvacation days
- Summarization for all days (vacation and nonvacation)

Complete the following steps to enable shifts, vacations, or both:

- Select when the beginning of the week starts.
- To configure shifts:
 - a. Select **Yes** in the **Specify shifts** drop-down list.
 - b. Optionally change the default settings for peak and off peak hours by selecting hours in the **Select Peak Hours** box.

Note: Changing the shift information after data has been summarized creates an inconsistency in the data. Data that was previously collected is not summarized again to account for the new shift values.

- To configure vacation settings:
 - a. Select **Yes** in the **Specify vacation days** drop-down list to enable vacation days.
 - b. Select **Yes** in the drop-down list if you want to specify weekends as vacation days.
 - c. Select **Add** to add vacation days.
 - d. Select the vacation days you want to add from the calendar.

The days you select are displayed in the list box.

If you want to delete any days you have previously chosen, select them and click **Delete**.

Note:

- 1) Add vacation days in the future. Adding vacation days in the past creates an inconsistency in the data. Data that was previously collected is not summarized again to account for vacation days.
- 2) Enabling shifts or vacation periods can significantly increase the size of the warehouse database. It will also negatively affect the performance of the Summarization and Pruning Agent.

32. Select the **Log Settings** check box to set the intervals for log pruning:




Figure 29. Log Settings pane of Summarization and Pruning Agent configuration window

- Select Prune WAREHOUSELOG, select the number of units for which data should be kept, and the unit of time (day, month or year).
- Select Prune WAREHOUSEAGGREGLOG, select the number of units for which data should be kept, and the unit of time (day, month or year).

33. Specify additional summarization and pruning settings in the **Additional Settings** pane:

The screenshot shows a window titled "Configure Warehouse Summarization and Pruning Agent". On the left is a tree view with the following items: Database Type, Sources, Scheduling, Work Days, Log Settings, and Additional Settings (which is selected). The main area is titled "Additional Settings Details" and contains the following settings:

- *Number of worker threads: 2
- *Maximum rows per database transaction: 1000
- *Use timezone offset from: Agent (dropdown)
- *Aggregate hourly data older than: 1
- *Aggregate daily data older than: 0
- *Maximum number of node errors to display: 100
- *Maximum number of Summarization and Pruning runs to display: 100
- *Database Connectivity Cache Time (minutes): 10
- *Batch mode: Single System (dropdown)
- *Database compression: Yes (dropdown)
- *Database Table Partitioning: Yes (dropdown)
- *Number of future partitions to maintain: 10
- *Group by threshold (days): 0
- Default table container: (empty text field)
- Default index container: (empty text field)

At the bottom of the window are five buttons: Back, Next, Home, OK, and Cancel.

Figure 30. Additional Settings pane of Summarization and Pruning Agent configuration window

- Specify the number of additional threads you want to use for handling summarization and pruning processing. The number of threads should be $2 * N$, where N is the number of processors running the Summarization and Pruning Agent. A higher number of threads can be used, depending on your database configuration and hardware.
- Specify the maximum rows that can be deleted in a single pruning transaction. Any positive integer is valid. The default value is 1000. There is no value that indicates you want all rows deleted.

If you increase the number of threads, you might consider increasing this value if your transaction log allows for it. The effective number of rows deleted per transaction is based on this value divided by the number of worker threads.

- Indicate a time zone for historical data from the **Use timezone offset from** drop down list. This field indicates which time zone to use when a user specifies a time period in a query for monitoring data.
 - Select **Agent** to use the time zone (or time zones) where the monitoring agents are located.

- Select **Warehouse** to use the time zone where the Summarization and Pruning Agent is located. If the Tivoli Data Warehouse and the Summarization and Pruning Agent are in different time zones, the **Warehouse** choice indicates the time zone of the Summarization and Pruning Agent, not the warehouse.

Skip this field if the Summarization and Pruning Agent and the monitoring agents that collect data are all in the same time zone.

- d. Specify the age of the data you want summarized in the **Aggregate hourly data older than** and **Aggregate daily data older than** fields. The default value is 1 for hourly data and 0 for daily data.
- e. The **Maximum number of node errors to display** refers to the node error table in the Summarization and Pruning workspace. It determines the maximum number of rows that workspace is to save and display.

- f. The **Maximum number of summarization and pruning runs to display** refers to the Summarization and Pruning Run table in the Summarization and Pruning workspace. It determines the maximum number of rows that workspace is to save and display.

Maximum number of Summarization and Pruning runs to display and Maximum number of node errors to display together determine the number of rows shown in the Summarization and Pruning overall run table and Errors table respectively. There is a minimum value of 10 for each. These equate to keywords KSY_SUMMARIZATION_UNITS and KSY_NODE_ERROR_UNITS in file KSYENV/sy.ini.

- g. The **Database Connectivity Cache Time** determines how long after a positive check for connectivity that the result will be cached. Longer times may result in inaccurate results in the workspace; however, it saves processing time.

Database Connectivity Cache Time records the number of minutes to cache the database connectivity for reporting purposes. The minimum value is 5 minutes. This equates to keyword KSY_CACHE_MINS in file KSYENV/sy.ini.

- h. **Batch mode** determines if data from different managed systems are used in the same database batch; this setting also improves performance.

Batch mode controls the batching method used by the Summarization and Pruning Agent. A value of Single Managed System (0) means that data should only be batched for the same system, whereas a value of Multiple Managed System (1) means that data from multiple systems can be batched together; this can lead to higher performance at potentially bigger transaction sizes. The default value is Single Managed System (0). This equates to keyword KSY_BATCH_MODE in file KSYENV/sy.ini.

- i. Specify whether you want to turn **Database compression** on or off.
- j. Specify whether you want to turn **Database Table Partitioning** on or off. When partitioning is enabled, the Summarization and Pruning Agent performs the following functions:
 - Creates new summary tables as partitioned tables.
 - Adds partitions to existing partitioned tables (including the detail tables created by the warehouse proxy) based on your configuration.
 - Prunes data from partitioned tables by dropping the appropriate partitions.

To change these values, you can either use the Summarization and Pruning configuration window's **Additional settings** tab or update these parameters directly in file KSYENV/sy.ini.

34. Save your settings and close the window. Click **Save** to save your settings. Click **Close** to close the configuration window.
35. Click **Finish** to complete the installation.

Results

When the installation is complete, the Manage Tivoli Enterprise Monitoring Services window is displayed. This window is used for starting, stopping, and configuring installed components. This

window can be displayed at any time from the Windows Start menu by clicking **Start → Programs → IBM Tivoli Monitoring → Manage Tivoli Monitoring Services**.

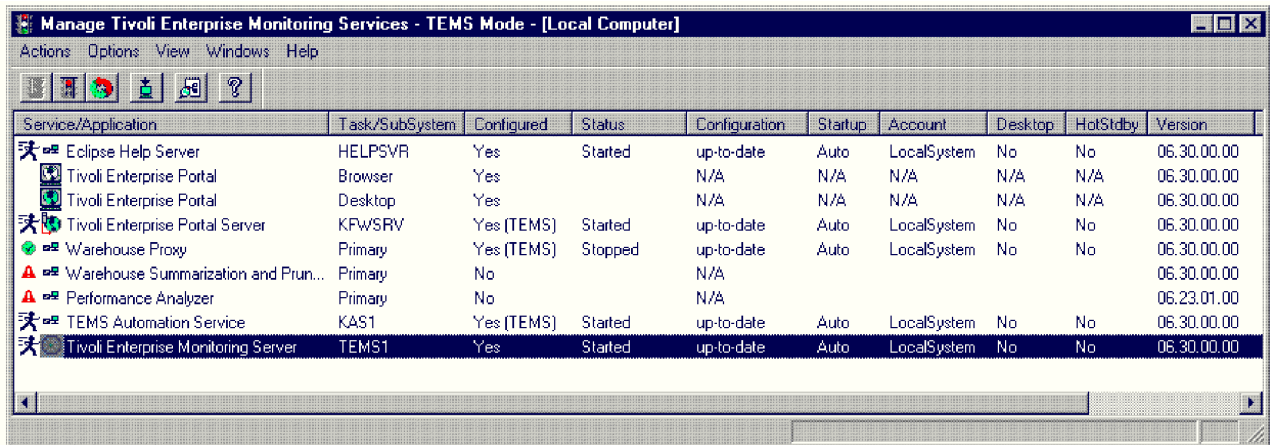


Figure 31. Manage Tivoli Enterprise Monitoring Services window

The symbol next to a component indicates its current state:

- A blue running figure indicates the component is started.
- A green check mark indicates that the component is configured and can be started.
- A red exclamation mark indicates that the component needs to be configured before it can be started.

What to do next

Complete the following tasks after you finish the installation procedure:

- Start the Eclipse Help Server and the Tivoli Enterprise Portal Server, if not started. Right-click the component in the Manage Tivoli Enterprise Monitoring Services window and select **Start**.
To log on to the Tivoli Enterprise Portal using the desktop client, double-click the Tivoli Enterprise Portal icon on your Windows desktop. Use the **sysadmin** password that you specified in Step 9 on page 173. (For more information about logging on to the Tivoli Enterprise Portal, see “Starting the Tivoli Enterprise Portal client” on page 366.)
- Configure the Summarization and Pruning Agent.
The installation procedure automatically starts and configures all the monitoring agents that you installed except for the Summarization and Pruning Agent. The Summarization and Pruning Agent is not configured and started during installation to give you an opportunity to configure history collection in advance for all installed monitoring agents, a task that must be performed prior to starting the Summarization and Pruning Agent for the first time.
To configure and start the Summarization and Pruning Agent, see the following procedures:
 - “Configuring the Summarization and Pruning Agent (JDBC connection)” on page 643
 - “Starting the Summarization and Pruning Agent” on page 656
- If you are using Microsoft SQL Server for the Tivoli Data Warehouse database, perform the following additional steps:
 - Create a schema with the same name (and owner) as the database user ID for accessing the Tivoli Data Warehouse. (The default user ID is ITMUser.) Change the default schema from dbo to this database user ID.
You specified the database user ID on the Configure SQL Data Source for Warehouse Proxy window during installation. See Step 12l of the installation procedure.
 - Ensure that the database is set up to support inbound network TCP/IP connections.

- Enable authentication of user credentials (optional)

During the installation and configuration of the monitoring server and portal server, user validation was not enabled so any user can use the Tivoli Enterprise Portal client, run tacmd commands, or issue SOAP server requests. If you want to enable user authentication through either the portal server (for LDAP authentication and single sign-on capability) or the monitoring server (for LDAP or local operating system authentication). See *Enabling user authentication* in the *IBM Tivoli Monitoring Administrator's Guide*.

- Install and setup a dashboard environment (optional)

If you want to install IBM Dashboard Application Services Hub with monitoring dashboard applications or custom dashboards, see “Required software and memory requirements for a dashboard environment” on page 91 and review the hardware and software requirements for Jazz for Service Management to determine if it can be installed on the same computer as IBM Tivoli Monitoring, or if you need to install it on a separate computer. After installing IBM Dashboard Application Services Hub, see “Installing and configuring the IBM Infrastructure Management Dashboards for Servers” on page 306 if you plan to use that dashboard application. Also, if you enabled LDAP user authentication through the portal server and you plan to use authorization policies to control access to monitored resources in the dashboards, see “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318.

Chapter 10. Installing IBM Tivoli Monitoring

You install the Tivoli Monitoring distributed components using the following DVDs or DVD images.

1. Base DVDs or DVD images contains:

- The IBM Tivoli Monitoring server and portal client components
- Warehouse Proxy Agent, Summarization and Pruning Agent, and Tivoli Performance Analyzer Agent and their application support
- Operating system agent application support

Note: There are operating system specific versions of the Base DVD.

2. Multiple agent DVDs or DVD images are provided for the following sets of monitoring agents included with IBM Tivoli Monitoring:

- The operating system agents (including the agentless OS monitoring agents) are on the *IBM Tivoli Monitoring Agents* DVD or DVD image
- System P agents are on the *Tivoli Monitoring Agents for System p*® DVDs or DVD image. There is a separate DVD with the application support for these agents.
- The Log File Monitoring agent is on the *IBM Tivoli Log File Agent* DVD or DVD image.

Note: Other products provide their own monitoring agent installation media.

3. The *Dashboards for Servers and Authorization Policy Components* DVD and DVD image contains the IBM Tivoli Monitoring components that are installed and used with the IBM Dashboard Application Services Hub component of Jazz for Service Management.

The installation and initial configuration of your environment consists of the steps described in Table 33 on page 194.

- This chapter contains the instructions to perform installation and initial configuration using a graphical user interface or a command line interface. If you want to perform a silent installation, see the instructions in Appendix B, “Performing a silent installation of IBM Tivoli Monitoring,” on page 841.
- The installation procedures provide information for installing a single component (such as the monitoring server) on one computer. If you want to install multiple components (such as the monitoring server and the portal server) on the same computer and you want to install them simultaneously, the actual steps might vary. See the individual sections for required configuration information during your installation.
- The installer validates configuration values for hub, remote, and standby monitoring servers. During configuration, the installer may prompt you to verify values that it considers suspect.
- The installation procedures in this chapter contain instructions for new installations. Follow the same procedures for upgrading or updating your installation. An *upgrade* is an installation that replaces a previous release or fix pack level of the product or component with a later release or fix pack level. An *update* is a modification to an existing installation at the same release or fix pack level.

For more information on fix pack installation, see “Installing product maintenance” on page 370.

- If your Tivoli Enterprise Portal Server is configured to a DB2 database, and you are upgrading your installation from a previous release, DB2 must be started before you choose to upgrade your Tivoli Enterprise Portal Server.
- See Appendix A, “Installation worksheets,” on page 833 for a set of worksheets that you can use to gather the information required for installation.
- The following sections contain Windows, Linux, and UNIX procedures for installing the various components. Use the procedure that best applies to your environment layout. For example, you can install a monitoring server on UNIX, a portal server on Linux, and a portal desktop client on Windows.

To install a monitoring server or monitoring agents on a z/OS system, see the appropriate z/OS documentation as described in Table 33.

- If your site uses the IP.PIPE or IP.SPIPE protocols for communications, be aware of the port usage rules for Tivoli Monitoring components. See “Controlling port number assignments” on page 421 for more information.

High-level installation steps

This section outlines the high-level installation steps and where you can find more information.

Table 33. IBM Tivoli Monitoring high-level installation steps

Step	Where to find information
Install the hub Tivoli Enterprise Monitoring Server.	<ul style="list-style-type: none"> • To install a hub monitoring server on a Windows, Linux, or UNIX system, see “Installing and configuring the hub Tivoli Enterprise Monitoring Server” on page 196. • To install a hub monitoring server on a z/OS system, see the IBM Tivoli Monitoring publication entitled <i>Configuring Tivoli Enterprise Monitoring Server on z/OS</i>.
Install any remote monitoring servers.	<ul style="list-style-type: none"> • To install a remote monitoring server on a Windows, Linux, or UNIX system, see “Installing and configuring the remote monitoring servers” on page 211. • To install a remote monitoring server on a z/OS system, see the IBM Tivoli Monitoring publication entitled <i>Configuring Tivoli Enterprise Monitoring Server on z/OS</i>.
Install the monitoring automation server if you plan to use its OSLC Performance Monitoring service provider component to integrate with the Jazz for Service Management Registry Services component and other products using OSLC linked data interfaces.	<ul style="list-style-type: none"> • For details on prerequisite software that should be installed before installing and configuring the monitoring automation server, see “Required software and memory requirements for using the OSLC Performance Monitoring service provider” on page 94. • For details on installing and configuring the automation server, see “Installing and configuring Tivoli Enterprise Monitoring Automation Server” on page 219.
Install the Tivoli Enterprise Portal Server.	“Installing the Tivoli Enterprise Portal Server” on page 226
Install monitoring agents.	<ul style="list-style-type: none"> • To install monitoring agents on Windows, Linux, or UNIX systems, see “Installing monitoring agents” on page 261. • To install monitoring agents on a z/OS system, see the configuration guides for your z/OS agent product. <p>If you plan to use the remote deployment function in IBM Tivoli Monitoring to install distributed monitoring agents across your environment, see Chapter 11, “Deploying monitoring agents across your environment,” on page 371 for the required steps. (You cannot use the remote deployment function to install z/OS monitoring agents.)</p> <p>If you used self-describing agents to install your product support, see “Self-describing agent installation” on page 394 to ensure that your environment is prepared for self-describing agents.</p>
Install non-agent bundles (optional).	“Working with non-agent bundles” on page 392
Install the portal desktop client (optional). If you want to use only the portal browser client or the portal Java WebStart client, you do not need to install the desktop client.	“Installing the Tivoli Enterprise Portal desktop client” on page 271

Table 33. IBM Tivoli Monitoring high-level installation steps (continued)

Step	Where to find information
Install required application support on the monitoring server, portal server, and portal desktop client.	"Installing and enabling application support" on page 275
Install and configure the IBM Infrastructure Management Dashboards for Servers with the IBM Dashboard Application Services Hub component of Jazz for Service Management if you are setting up a monitoring dashboard environment.	<ul style="list-style-type: none"> For details on prerequisite software that should be installed before installing and configuring the dashboards for servers, see "Required software and memory requirements for a dashboard environment" on page 91. For information on the installation options for installing the dashboards for servers component, see "Using IBM Installation Manager to install or update IBM Tivoli Monitoring components for a dashboard environment" on page 299. For details on installing and configuring dashboards for servers, see "Installing and configuring the IBM Infrastructure Management Dashboards for Servers" on page 306.
Install and configure the Tivoli Authorization Policy Server and tivcmd Command Line Interface for Authorization Policy with the IBM Dashboard Application Services Hub component of Jazz for Service Management if you are setting up a monitoring dashboard environment and plan to use authorization policies to control what managed systems and managed system groups users can access from the dashboards.	<ul style="list-style-type: none"> For details on prerequisite software that should be installed before installing and configuring the authorization policy server, see "Required software and memory requirements for a dashboard environment" on page 91. For information on the installation options for installing the authorization policy server and tivcmd CLI, see "Using IBM Installation Manager to install or update IBM Tivoli Monitoring components for a dashboard environment" on page 299. For details on installing and configuring the authorization policy server and tivcmd CLI, see "Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy" on page 318.
Install domain definitions for Tivoli Performance Analyzer if you are using the Tivoli Performance Analyzer monitoring agent.	"Installing domain definitions for Tivoli Performance Analyzer" on page 336
Install the language packs for all languages other than English.	"Installing language packs" on page 339
Configure the clients, browser and Java runtime environment.	<ul style="list-style-type: none"> "Configuring clients, browsers, and JREs" on page 342 "Installing and configuring IBM Java 7" on page 348
Specify what browser to use to display the online help.	"Specifying the browser used for online help" on page 363
Start the Tivoli Enterprise Portal to verify your installation.	"Starting the Tivoli Enterprise Portal client" on page 366
Use Web Start to download the run the desktop client.	"Using Web Start to download and run the desktop client" on page 367
Enable user authentication and authorization for the Tivoli Enterprise Portal.	See the <i>IBM Tivoli Monitoring Administrator's Guide</i>
Set up your dashboard environment which includes configuring a data provider connection from the Dashboard Application Services Hub to the portal server and enabling user authentication and authorization.	See <i>Preparing your Dashboard Environment</i> in the <i>IBM Tivoli Monitoring Administrator's Guide</i> .

Installing and configuring the hub Tivoli Enterprise Monitoring Server

The following sections provide detailed information for installing and initially configuring the hub monitoring server.

- “Windows: Installing the hub monitoring server”
- “Linux or UNIX: Installing the hub monitoring server” on page 200

Note: IBM Tivoli Monitoring does not support multiple hub monitoring servers on a single Windows, Linux, or UNIX computer or LPAR.

Windows: Installing the hub monitoring server

Follow the steps in this section to install the hub monitoring server on a Windows computer.

Procedure

1. Launch the installation wizard by double-clicking the setup.exe file on the IBM Tivoli Monitoring Base DVD or DVD image.

Note: If you are running Windows 2003 or Windows XP and have security set to check the software publisher of applications, you might receive an error stating that the setup.exe file is from an unknown publisher. Click **Run** to disregard this error message.

2. Click **Next** on the Welcome window.

Note: If you have another IBM Tivoli Monitoring component already installed on this computer, select **Modify** on the Welcome window to indicate that you are updating an existing installation. Click **OK** on the message telling you about preselected items. Then skip to Step 6.

3. Click **Accept** to accept the license agreement.
4. Choose the directory where you want to install the product. The default directory is C:\IBM\ITM. Click **Next**.
5. Type a 32-character encryption key. You can use the default key.

Note:

- a. Do not use any of the following characters in your key:

&	ampersand
	pipe
'	single quote
=	equal sign
\$	dollar sign

In addition, do not specify double-byte (DBCS) characters.

- b. Ensure that you document the value you use for the key. Use this key during the installation of any components that communicate with this monitoring server.

Click **Next** and then click **OK** to confirm the encryption key.

6. On the Select Features window, select the check box for **Tivoli Enterprise Monitoring Server**.

When you select the **Tivoli Enterprise Monitoring Server** check box, all of the check boxes in the attached subtree are automatically selected. The support check boxes in the subtree are for installing application support files for monitoring agents to the monitoring server. You can leave all of the support check boxes selected so you do not need to reconfigure application support as new agent types are added to your environment, but adding support for many agents can extend the installation time considerably. If you do not intend to install a particular type of agent, remove the check mark from the selection box. For detailed information about application support, see “Installing and enabling application support” on page 275.

For additional information about these components, press the **Help** button.

Note:

- a. If you have purchased monitoring agents that run on z/OS, but have not purchased IBM Tivoli Monitoring as a separate product, expand the **Tivoli Enterprise Monitoring Server** node. Clear the check boxes in the subtree for the base agents except for the Warehouse Proxy and Summarization and Pruning Agents. (The *base* monitoring agents are included with the base IBM Tivoli Monitoring installation package: see “Installing monitoring agents” on page 261.)
 - b. If you are updating an existing installation (you selected **Modify** on the Welcome window), all check boxes on the Select Features window reflect your choices during the initial installation. Clearing a check box has the effect of *uninstalling* the component. Clear a check box only if you want to remove a component.
 - c. If you have decided to use self-describing agents to install product application support, you do not have to select the support check box for this product.
7. Click **Next** to display the Agent Deployment window.

The Agent Deployment window lists monitoring agents on this installation image that you can add to the agent depot. The agent depot contains agents that you can deploy to remote computers. For information about how to deploy agents in the agent depot to remote computers, see Chapter 11, “Deploying monitoring agents across your environment,” on page 371.

For additional information about agent deployment, press the **Help** button.

Note: By default, the agent depot is located in the *itm_installdir/CMS/depot* directory on Windows. If you want to use a different directory, create the directory (if it does not exist) and specify the directory using the DEPOTHOME key in the KBBENV file.

Select the agents, if any, that you want to add to the agent depot. (You can add agents to the agent depot at a later time by updating your installation.) Click **Next**.

8. If no IBM Tivoli Monitoring component has been previously installed on this computer, a window is displayed for you to select a program folder for the Windows Start menu. Select a program folder, and click **Next**. The default program folder name is IBM Tivoli Monitoring.
9. If the TEPS Desktop and Browser Signon ID and Password window is displayed, enter and confirm the password to be used for logging on to the Tivoli Enterprise Portal. The default logon user ID, **sysadmin**, cannot be changed on this window.

This password is required only when **Security: Validate Users** is enabled on the hub monitoring server. This window is not displayed if the **sysadmin** user ID has already been defined in the operating system.
10. Review the installation summary details. The summary identifies the components you are installing. Click **Next** to begin the installation.

After the components are installed, a configuration window (called the Setup Type window) is displayed.
11. Clear the check boxes for any components that are already installed and configured (at the current release level) on this computer, unless you want to modify the configuration. Click **Next** to start configuring all selected components.
12. Configure the Tivoli Enterprise Monitoring Server:

- a. Select the type of monitoring server you are configuring: **Hub** or **Remote**. For this procedure, select **Hub**.
- b. Verify that the name of this monitoring server is correct in the **TEMS Name** field. If it is not, change it.

The default name is *HUB_host_name*, for example, *HUB_itmserv16*. *This name must be unique in the enterprise.*

- c. Identify the communications protocol for the monitoring server. You have four choices: IP.UDP, IP.PIPE, IP.SPIPE, or SNA. You can specify up to three methods for communication. If the method you identify as Protocol 1 fails, Protocol 2 is used as a backup. If Protocol 2 fails, Protocol 3 is used as a backup.

- d. Leaving **TEMS Security Compatibility Mode** enabled performs the following functions:
 - 1) It allows for backwards compatibility with previous versions of IBM Tivoli Monitoring components shipped before V6.3 when securing command execution (Situation Take Action, Execute Command, Take Action, Workflow Policy Take Action). Accept the default if you have a previous version of a portal server, hub monitoring server, or OMEGAMON product connecting to this Tivoli Enterprise Monitoring Server or if the system clocks are not synchronized across the systems with monitoring servers, the portal server, and monitoring agents.
 - 2) It permits commands to still execute even if cryptographic facilities are not available. This gives you time to ensure that, for example, your z/OS Tivoli Enterprise Monitoring Server has the Integrated Cryptographic Services Facility (ICSF) available.

Note: If you disable this option and you have a mixture of pre-V6.3 and V6.3 server components, if clocks are not synchronized, or if you don't have cryptographic services available, then command executions might fail.

- e. Do not select any of the other options on this window (for example **Address Translation**, **Tivoli Event Integration Facility** or the option to configure **Hot Standby**). You can configure these options after installation is complete.
- f. Remove the check mark from the box for **Security: Validate Users**. Use the procedures in the *IBM Tivoli Monitoring Administrator's Guide* to configure security after installation is complete. If you decide to leave security enabled, and you want to use LDAP to authenticate users instead of the hub security system, use the information in the *IBM Tivoli Monitoring Administrator's Guide* to complete the configuration.
- g. Click **OK**.
- h. Complete the following fields for the communications protocol for the monitoring server.

Table 34. Communications protocol settings for the hub monitoring server

Field	Description
IP.UDP Settings	
Hostname or IP Address	The host name or IP address for the hub monitoring server.
Port # or Port Pools	The listening port for the hub monitoring server. The default port is 1918.
IP.PIPE Settings	
Hostname or IP Address	The host name or IP address for the hub monitoring server.
Port Number	The listening port for the monitoring server. The default number is 1918.
IP.SPIPE Settings	
Hostname or IP Address	The host name or IP address for the hub monitoring server.
Port number	The listening port for the hub monitoring server. The default value is 3660.
SNA Settings	
Network Name	The SNA network identifier for your location.
LU Name	The LU name for the monitoring server. This LU name corresponds to the Local LU Alias in your SNA communications software.
LU 6.2 LOGMODE	The name of the LU6.2 LOGMODE. The default value is CANCTDCS.

Table 34. Communications protocol settings for the hub monitoring server (continued)

Field	Description
TP Name	The transaction program name for the monitoring server.

- i. If you are certain that you have typed the values for all of these fields with *exactly* the correct casing (upper and lower cases), you can select **Use case as typed**. However, because IBM Tivoli Monitoring is case-sensitive, consider selecting **Convert to upper case** to reduce the chance of user error.
- j. If you are installing a Tivoli Enterprise Monitoring Automation Server together with the Tivoli Enterprise Monitoring Server, some additional configuration steps are required. See step 14 on page 221 in “Windows: Installing a Tivoli Enterprise Monitoring Automation Server” on page 220.
- k. Click **OK** to continue.

Press the **Help** button for additional information about the configuration parameters of the monitoring server.

13. If you selected **Tivoli Event Integration Facility**, provide the host name and port number for the Netcool/OMNIBus EIF probe or TEC event server to which you want to forward events and click **OK**.

The default port number for the Netcool/OMNIBus EIF Probe Version 10 or later is 9998. The default port number for the TEC event server is 5529. However, if you used Tivoli Business Service Manager V4.2.1 to install the EIF probe then the probe's default port number is 5530.

Note: Before configuring the Tivoli Event Integration Facility support, verify that you have performed other prerequisite steps to integrate with the event server. See Chapter 31, “Setting up event forwarding to Tivoli Enterprise Console,” on page 799 or Chapter 30, “Setting up event forwarding to Netcool/OMNIBus,” on page 697 for details.

14. Enable application support on the monitoring server.

In Step 6 on page 196, you selected the base monitoring agents for which you wanted to install application support files on the monitoring server. In this step, you activate the application support through a process known as *seeding* the monitoring server.

Note:

- If you are running in a Hot Standby environment, shut down your Hot Standby (that is, mirror) monitoring server before completing this procedure. You may restart the Hot Standby monitoring server only after you have seeded the hub server.
- If you are using self-describing agent capability to install monitoring server support, you might be able to skip this step. This depends on the method of seeding you want to use for a given self-describing agent product. If automatic self-describing agent seeding is required, then you can skip this step. If you want to manually seed the self-describing agent installed product, then you must perform this step. For more information, see “Configuring self-describing agent seeding” on page 206.

- a. Specify the location of the monitoring server to which to add application support. You have two choices:

- **On this computer**
- **On a different computer**

Click **OK**.

For additional information about these parameters, press the **Help** button.

- b. Click **OK** on the Select the application support to add to the TEMS window.

This window lists the monitoring agents that you selected in Step 6 on page 196. Click **OK** to begin seeding the monitoring server (using the SQL files listed on this window).

This process can take up to 20 minutes. As the seeding process completes, a progress bar is displayed, showing the progress of seeding. In turn, the application support for the agents you

selected.




Figure 32. Progress bar for application seeding

Once seeding completes, if support could not be added, a window is displayed showing all seeding results.

- c. Click **Next** on the message that provides results for the process of adding application support (see Figure 54 on page 284). A return code of 0 (rc: 0) indicates that the process succeeded.

Note: If the Application Support Addition Complete window is not displayed after 20 minutes, look in the IBM\ITM\CNPS\Logs\seedkpc.log files (where *pc* is the two-character product code for each monitoring agent) for diagnostic messages that help you determine the cause of the problem. For a list of product codes, see Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883.

15. Configure the communication between any IBM Tivoli Monitoring component and the hub monitoring server:
 - a. Specify the default values for IBM Tivoli Monitoring components to use when they communicate with the monitoring server.
 - 1) If agents must cross a firewall to access the monitoring server, select **Connection must pass through firewall**.
 - 2) Identify the type of protocol that the agents use to communicate with the hub monitoring server. You have four choices: IP.UDP, IP.PIPE, IP.SPIPE, or SNA. You can specify up to three methods for communication. If the method you identify as Protocol 1 fails, Protocol 2 is used as a backup. If Protocol 2 fails, Protocol 3 is used as a backup.

Click **OK**.

- b. Complete the communication protocol fields for the monitoring server. See Table 34 on page 198 for definitions of these fields. Click **OK**.

For additional information about these parameters, press the **Help** button.

16. Click **Finish** to complete the installation.
17. Click **Finish** on the Maintenance Complete window if you are updating an existing installation.

Results

The Manage Tivoli Enterprise Monitoring Services utility is opened. (This might take a few minutes.) You can start, stop, and configure IBM Tivoli Monitoring components with this utility.

Linux or UNIX: Installing the hub monitoring server

Follow the steps in this section to install and configure the hub monitoring server on a Linux or UNIX computer.

Table 35. Steps for installing a hub monitoring server on a Linux or UNIX computer

Steps	Where to find information
Install the hub monitoring server.	“Installing the monitoring server” on page 201
Configure the hub monitoring server.	“Configuring the hub monitoring server” on page 202

Table 35. Steps for installing a hub monitoring server on a Linux or UNIX computer (continued)

Steps	Where to find information
Add application support to the hub monitoring server.	"Adding application support to the hub monitoring server" on page 208

Installing the monitoring server

Follow the steps in this section to install the monitoring server on a Linux or UNIX computer.

Procedure

1. In the directory where you extracted the installation files, run the following command:
`./install.sh`
2. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default (/opt/IBM/ITM). If you want to use a different installation directory, type the full path to that directory and press Enter.

Note: You must not specify the path of the directory containing `./install.sh` as your IBM Tivoli Monitoring home directory. On certain platforms, this can cause the plug-in JAR files to overwrite themselves and become zero length files. The installation will fail as a result.

In addition, do not use the number sign (#) or exclamation mark (!) characters anywhere in your installation directory path.

3. If the directory you specified does not exist, you are asked whether to create it. Type 1 to create this directory.
4. The following prompt is displayed:
 Select one of the following:
 1) Install products to the local host.
 2) Install products to depot for remote deployment (requires TEMS).
 3) Install TEMS support for remote seeding
 4) Exit install.
 Please enter a valid number:

Type 1 to start the installation and press Enter.

The end user license agreement is displayed. Press Enter to read through the agreement.

5. Type 1 to accept the agreement and press Enter.
 If you are upgrading IBM Tivoli Monitoring from global GSKit enabled to local GSKit enabled, the following text is displayed:
 Local GSKit has been installed for IBM Tivoli Monitoring. The global GSKit installation will no longer be used by IBM Tivoli Monitoring. It will not be uninstalled automatically, because it might be used by other applications. You may uninstall it manually, if IBM Tivoli Monitoring is the only application utilizing it.
6. Type a 32-character encryption key and press Enter. If you want to use the default key, press Enter without typing any characters.

Note:

- a. Do not use any of the following characters in your key:
 & ampersand
 | pipe
 ' single quote
 = equal sign
 \$ dollar sign

In addition, do not specify double-byte (DBCS) characters.

- b. Ensure that you document the value you use for the key. Use this key during the installation of any components that communicate with this monitoring server.

Note: If you are upgrading IBM Tivoli Monitoring instead of a first-time installation, you will be prompted by the installer to upgrade the common prerequisite components for agents and servers. The product packages available for this operating system and component support categories are listed.

7. Type the number that corresponds to your operating environment. For a new installation, type **1** IBM Tivoli Monitoring components for this operating system. This option installs all components for your operating system. Options 2 - 4 install application support to the server components listed. Select **5** Other operating systems to install components for an operating system other than the one detected.
8. Type **4** to install the monitoring server for your current operating system. Press Enter.
A list of the components to install is displayed.
9. Type **1** to confirm the installation.
The installation begins.
10. For first-time installations only, when prompted type a name for your monitoring server. For example: `HUB_hostname`. Do not use the fully qualified host name. Press Enter.

Note: When installing the monitoring server, the application support available on the installation media for the monitoring server is installed automatically.

11. When you finish installing the monitoring server, you are asked if you want to add application support to it.

Note: If you are using self-describing agent capability to install monitoring server support, you might be able to skip this step. This depends on the method of seeding you want to use for a given self-describing agent product. If automatic self-describing agent seeding is required, then you can skip this step. If you want to manually seed the self-describing agent installed product, then you must perform this step. For more information, see “Configuring self-describing agent seeding” on page 206.

12. After all of the components are installed, you are asked whether you want to install components for a different operating system. Type **2** and press Enter.
13. Installation is complete. If your IBM Tivoli Monitoring environment is not already secured you will be asked at this point if you want to secure it. If your IBM Tivoli Monitoring environment is already secured this question is skipped. The product installation process creates the majority of directories and files with world write permissions. IBM Tivoli Monitoring provides the `secureMain` utility to help you keep the monitoring environment secure. You can secure your installation now, or manually execute the `secureMain` utility later. For more information, see Appendix G, “Securing your IBM Tivoli Monitoring installation on Linux or UNIX,” on page 925.

What to do next

The next step is to configure your monitoring server.

Configuring the hub monitoring server

To configure your hub monitoring server, complete the following steps.

Before you begin

If you will be configuring user security, you should use root login ID to configure.

Procedure

1. At the command-line, change to the /opt/IBM/ITM/bin directory (or the directory where you installed IBM Tivoli Monitoring).
2. Run the following command:

```
./itmcmd config -S -t tems_name
```

where *tems_name* is the name of your monitoring server (for example, HUB_itmdev17).

3. Press Enter to indicate that this is a hub monitoring server (indicated by the *LOCAL default).
4. Press Enter to accept the default host name for the monitoring server. This should be the host name for your computer. If it is not, type the correct host name and then press Enter.
5. Enter the type of protocol to use for communication with the monitoring server. You have four choices: ip.udp, ip.pipe, ip.spipe, or sna. Press Enter to use the default communications protocol (IP.PIPE).
6. If you want to set up a backup protocol, enter that protocol and press Enter. If you do not want to use backup protocol, press Enter without specifying a protocol.
7. Depending on the type of protocol you specified, provide the following information when prompted:

Table 36. UNIX monitoring server protocols and values

Protocol	Value	Definition
IP.UDP	IP Port Number	The port number for the monitoring server. The default is 1918.
SNA	Net Name	The SNA network identifier for your location.
	LU Name	The LU name for the monitoring server. This LU name corresponds to the Local LU Alias in your SNA communications software.
	Log Mode	The name of the LU6.2 LOGMODE. The default value is "CANCTDCS."
IP.PIPE	IP.PIPE Port Number	The port number for the monitoring server. The default is 1918.
IP.SPIPE	IP.SPIPE Port Number	The port number for the monitoring server. The default is 3660.

8. Press Enter to *not* specify the name of the KDC_PARTITION. You can configure the partition file at a later time, as described in Appendix C, "Firewalls," on page 861.
9. Press Enter when prompted for the path and name of the KDC_PARTITION.
10. If you want to use Configuration Auditing, press Enter. If you do not want to use this feature, type 2 and press Enter.
11. Press Enter to accept the default setting for the Hot Standby feature (none).
For best results, wait until after you have fully deployed your environment to configure the Hot Standby feature for your monitoring server. See the *IBM Tivoli Monitoring High Availability Guide for Distributed Systems* for information about configuring Hot Standby.
12. Press Enter to accept the default value for the Optional Primary Network Name (none).
13. Press Enter for the default **Security: Validate User** setting (NO).

Important: Do *not* change this to set **Security: Validate User**. You can configure security after configuring the monitoring server, as described in the *IBM Tivoli Monitoring Administrator's Guide*.

14. If you want to forward situation events to either IBM Tivoli Netcool/OMNIBus console or the IBM Tivoli Enterprise Console (TEC), type 1 and press Enter to enable the Tivoli Event Integration Facility. Complete the following additional steps:
 - a. For EIF Server, type the hostname of the TEC event server or the hostname of the Netcool/OMNIBus EIF probe and press Enter.
 - b. For EIF Port, type the EIF reception port number for the TEC event server or the Netcool/OMNIBus EIF probe and press Enter.

The default port number for the Netcool/OMNIBus EIF Probe Version 10 or later is 9998. The default port number for the Tivoli Enterprise Console event server is 5529. However, if you used Tivoli Business Service Manager V4.2.1 to install the EIF probe then the probe's default port number is 5530.

Note: Before configuring the Tivoli Event Integration Facility support, verify that you have performed other prerequisite steps to integrate with the event server. See Chapter 31, “Setting up event forwarding to Tivoli Enterprise Console,” on page 799 or Chapter 30, “Setting up event forwarding to Netcool/OMNIBus,” on page 697 for details.

15. To disable Workflow Policy Activity event forwarding, type **1** and press Enter. Otherwise, press Enter to accept the default value (2=NO). See the note in “Event integration with Tivoli Enterprise Console” on page 800 for more information.
16. To disable TEMS Security Compatibility Mode, type **2** and press Enter. Otherwise, press Enter to accept the default value (1=YES).

Leaving **TEMS Security Compatibility Mode** enabled performs the following functions:

- a. It allows for backwards compatibility with previous versions of IBM Tivoli Monitoring components shipped before V6.3 when securing command execution (Situation Take Action, Execute Command, Take Action, Workflow Policy Take Action). Accept the default if you have a previous version of a portal server, hub monitoring server, or OMEGAMON product connecting to this Tivoli Enterprise Monitoring Server or if the system clocks are not synchronized across the systems with monitoring servers, the portal server, and monitoring agents.
- b. It permits commands to still execute even if cryptographic facilities are not available. This gives you time to ensure that, for example, your z/OS Tivoli Enterprise Monitoring Server has the Integrated Cryptographic Services Facility (ICSF) available.

Note: If you disable this option and you have a mixture of pre-V6.3 and V6.3 server components, if clocks are not synchronized, or if you don't have cryptographic services available, then command executions might fail.

17. Type **6** to accept the default SOAP configuration and exit the configuration.

Note: You can configure any SOAP information at a later time. See Chapter 17, “Configuring IBM Tivoli Monitoring Web Services (the SOAP Server),” on page 461 for information.

Results

The monitoring server is now configured.

A configuration file is generated in the *install_dir/config* directory with the format *host_name_ms_tems_name.config* (for example, itmdev17_ms_HUBitmdev17.config).

What to do next

If you are installing a Tivoli Enterprise Monitoring Automation Server together with the Tivoli Enterprise Monitoring Server, some additional configuration steps are required. For more information, see “Configuring the Tivoli Enterprise Monitoring Automation Server” on page 224.

Enabling self-describing agent capability at the hub monitoring server

This feature integrates the installation of an agent with the dispersal and installation of associated product support files throughout your IBM Tivoli Monitoring infrastructure.

The self-describing agent feature makes it possible for new or updated Tivoli Monitoring agents to become operational after installation, without having to perform additional product support installation

steps. Self-describing agents apply version updates to other components automatically without the need to recycle your hub Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, or remote Tivoli Enterprise Monitoring Server.

Self-describing agent requirements

To take advantage of self-describing agent capability, your Tivoli Management Services server components must be at Version 6.2.3 or higher.

By default, the self-describing agent capability is disabled at the hub monitoring server. Self-describing agent environment variables are enabled by default for remote monitoring servers, portal servers, and self-describing agent enabled agents. However, these components disable self-describing agent capability if connected to a hub monitoring server that has the self-describing agent capability disabled.

Enabling the self-describing agent capability at the hub monitoring server controls the capability across all components. You can disable the capability individually at a remote monitoring server, portal server, or agent, as described in “Enabling or disabling self-describing agent capability” on page 396.

After installing and configuring the hub monitoring server, start the hub monitoring server if not already running in order to specify the default self-describing agent seeding options. Set the **KMS_SDA** environment variable to enable self-describing agent capability. Then recycle the hub monitoring server.

Beginning with V6.2.3 fix pack 1, self-describing agent application support installation is supported in Hot-Standby (FTO) configuration. The **KMS_SDA** settings should match on the primary and standby monitoring servers. An error message will report a discrepancy if the primary and standby monitoring servers have different **KMS_SDA** settings.

Dynamically controlling the hub monitoring server self-describing agent capability

When the hub Tivoli Enterprise Monitoring Server is configured for self-describing agent capability (by setting the variable **KMS_SDA=Y**) and started, you can dynamically change the operational state of the hub self-describing agent feature without recycling the hub monitoring server.

It is the hub monitoring server that controls the self-describing agent function for all remote monitoring servers and portal servers, provided these components are started and with the self-describing agent capability turned on. The following table describes the administration commands to dynamically control the self-describing agent function at the hub monitoring server:

Table 37. tacmd commands to control the operational state of the hub self-describing agent function

tacmd Command	Description
<code>tacmd suspendsda</code>	Disable (suspend) self-describing agent capability without recycling the hub monitoring server.
<code>tacmd resumesda</code>	Enable (resume) self-describing agent capability without recycling the hub monitoring server.
<code>tacmd listsdastatus</code>	Confirm if the self-describing agent capability is enabled or disabled for each hub or remote Tivoli Enterprise Monitoring Server.

For the full syntax, including parameter descriptions, see the *IBM Tivoli Monitoring Command Reference*.

You can issue the `tacmd suspendsda` command if you want to turn off the self-describing agent function during critical monitoring periods. Once the maintenance period arrives, you can issue the `tacmd resumesda` command to allow any available self-describing agents to upgrade server product support again, if needed. When the self-describing agent function is changed from disabled (suspend) to enabled (resume), the hub monitoring server ensures that all running remote monitoring servers and all running self-describing agent enabled agents are notified to retry their installation attempt. If the notified agent

product support is allowed to install via the self-describing agent capability, and this version of its product support files are not yet installed on the hub or remote monitoring server, the self-describing agent installation takes place.

Note:

1. The `tacmd` commands `suspend_sda` and `resume_sda` work only when the hub monitoring server is started and self-describing agent capability is turned on by setting the variable `KMS_SDA=Y`. These commands cannot turn on the self-describing agent function if the self-describing agent capability configured is not configured (that is, when `KMS_SDA=N` at the hub monitoring server). In this case, running the commands has no effect and a warning is displayed.
2. You cannot run the `suspend_sda` and `resume_sda` commands against any remote monitoring server or Hot Standby (FTO) mirror hub monitoring server. The self-describing agent enablement status of the remote monitoring server is controlled only by the existing configuration variable `KMS_SDA` on the hub monitoring server.

Complete the following steps to suspend and then resume the self-describing agent capability without recycling the hub monitoring server:

1. Enable the hub monitoring server for self-describing agent capability by setting the variable `KMS_SDA=Y` and start the hub monitoring server. The hub monitoring server is now enabled for self-describing agent capability. Any products currently configured to allow self-describing agents to install will take place if the agents are running.
2. Issue the command `tacmd suspend_sda` to dynamically turn OFF the self-describing agent function for the hub monitoring server. This command also turns off all self-describing agent capability for all remote monitoring servers and portal servers.
3. Issue the command `tacmd resume_sda` to dynamically turn ON the self-describing agent function for the hub monitoring server. This command also turns on all self-describing agent capability for all remote monitoring servers and portal servers.
4. Issue the command `tacmd list_sda_status` to view the self-describing agent enablement status of all Tivoli Enterprise Monitoring Servers in the enterprise and to confirm that the enablement is complete.

Related concepts:

“Managing your self-describing agent installations” on page 394

You can use `tacmd` commands to exercise more control over the self-describing agent products that are installed by the self-describing agent process.

Configuring self-describing agent seeding

Two types of installation scenarios are available: a pristine installation and an upgrade installation.

A product installation is considered a *pristine installation* if you have no existing monitoring configuration definitions, specifically situation definitions already configured in the Tivoli Enterprise Monitoring Server for that product. A product installation is considered an *upgrade installation* if you have configured monitoring definitions, specifically situation definitions configured for the self-describing agent of the installed product.

For both the pristine installation and upgrade installation, use the `tacmd editSdaOptions` command to configure how the self-describing agent feature seeds product definitions.

Configuring seeding for a pristine installation

Configure one of the following values to define how distribution definitions are applied during the pristine installation of the specified product.

NEW Use the configuration choice of **NEW** to add all product distribution definitions to the Tivoli Enterprise Monitoring Server during a pristine installation.

NONE

Use the configuration choice of **NONE** to prevent product distribution definitions from being added to the Tivoli Enterprise Monitoring Server during a pristine installation.

ALL Use the configuration choice of **ALL** to add all product distribution definitions to the Tivoli Enterprise Monitoring Server during a pristine installation.

DISABLE

Use the configuration choice of **DISABLE** to prevent all product monitoring definitions from being added to the Tivoli Enterprise Monitoring Server during a pristine installation.

Configuring seeding for an upgrade installation

Configure one of the following values to define how distribution definitions are applied during the upgrade installation of the product identified.

NEW Use the configuration choice of **NEW** to prevent any product distribution definitions from being added to the Tivoli Enterprise Monitoring Server during a upgrade installation.

NONE

Use the configuration choice of **NONE** to prevent any product distribution definitions from being added to the Tivoli Enterprise Monitoring Server during a upgrade installation.

ALL Use the configuration choice of **ALL** to add all product distribution definitions to the Tivoli Enterprise Monitoring Server during a upgrade installation.

DISABLE

Use the configuration choice of **DISABLE** to prevent all product monitoring definitions from being added to the Tivoli Enterprise Monitoring Server during an upgrade installation.

Default installation seeding values

If no seeding values are configured, the default seeding values are used. The default installation value for a pristine installation is **ALL**. The default installation value for an upgrade installation is **NEW**.

If only a single installation value is provided to the `editSdaOptions` command, the value for the unspecified option is set to the default value, that is, **ALL** for a pristine installation and **NEW** for an upgrade installation.

For more information on the commands in this section, see the *IBM Tivoli Monitoring Command Reference*.

Example of using the `tacmd listsdaoptions` command when no record is found

```
tacmd listsdaoptions
```

```
KUILS0101I: No records found on the server.
```

This means default SDA seeding will occur: **ALL** for pristine install, **NEW** for upgrade install

Example of using the `tacmd editsdaoptions` command to set default behaviors

```
tacmd editSdaOptions -t DEFAULT -o SEEDING=DISABLE
```

No configuration options were found for the specified type. The following options will be created:

```
PRODUCT INSTALL SEED UPGRADE SEED  
DEFAULT DISABLE DISABLE
```

```
Are you sure you want to update the selected options? Type Y for yes. Type N for no.  
Y
```

```
KUIES0150I: The selected SDA configuration options records were successfully updated.
```

Example of using the tacmd editsdaoptions to set options for a specific agent

```
tacmd editSdaOptions -t NT LZ UX -o INSTALL_SEED=ALL UPGRADE_SEED=NEW
```

No configuration options were found for the specified type. The following options will be created:

```
PRODUCT INSTALL SEED UPGRADE SEED
LZ ALL NEW
NT ALL NEW
UX ALL NEW
```

Are you sure you want to update the selected options? Type Y for yes. Type N for no.
Y

KUIES0150I: The selected SDA configuration options records were successfully updated.

Example of using the tacmd listsdaoptions to display existing records

```
tacmd listsdaoptions
PRODUCT INSTALL SEED UPGRADE SEED
DEFAULT DISABLE DISABLE
LZ ALL NEW
NT ALL NEW
UX ALL NEW
```

Complete the following steps to enable self-describing agent capability

- **Windows**
 1. In the Manage Tivoli Enterprise Monitoring Services application, right-click **Tivoli Enterprise Monitoring Server** and select **Advanced** → **Edit ENV file**. The component environment variable file is displayed.
 2. Change the variable KMS_SDA=N to **KMS_SDA=Y**.
 3. Save and close the file.
 4. Recycle the **Tivoli Enterprise Monitoring Server** to have your changes take effect.
- **Linux** **UNIX**
 1. Change to the <install_dir>/config directory and open the coordinating file for the monitoring server: <hostname>_ms_<tems_name>.config.
 2. Change the variable KMS_SDA=N to **KMS_SDA=Y**.
 3. Save and close the file.
 4. Recycle the **Tivoli Enterprise Monitoring Server** to have your changes take effect.

Adding application support to the hub monitoring server

All monitoring agents require that application support files be installed on the monitoring servers (hub and remote), portal server, and portal desktop clients in your environment. Application support files contain the information required for agent-specific workspaces, helps, predefined situations, and other data.

When you installed the monitoring server on Linux or UNIX, following the instructions in “Installing the monitoring server” on page 201, the application support files for all base monitoring agents and other supported agents were automatically installed on the monitoring server. This is different to installing the monitoring server on Windows, in that the Linux or UNIX installation always automatically installs the application support files for monitoring agents. This might affect self-describing agents and you should always check the monitoring server and portal server to ensure that self-describing agent products are installed as expected.

The *base* monitoring agents are the monitoring agents included on the IBM Tivoli Monitoring installation media. Now you must enable the application support for those agents. The process of enabling application support is also referred to as *activating* or *adding* application support, and in the case of the monitoring server, as *seeding* the monitoring server.

Follow the instructions in this section to add application support for agents to a monitoring server on Linux or UNIX. For detailed information about application support, and for instructions to install and enable application support for nonbase agents, see “Installing and enabling application support” on page 275 and “Configuring application support for nonbase monitoring agents” on page 279.

Note: If you are running in a Hot Standby environment, shut down your Hot Standby (that is, mirror) monitoring server before completing this procedure. You may restart the Hot Standby monitoring server only after you have seeded the hub server.

Use one of the following procedures to add application support for Base monitoring agents to a monitoring server:

- “Command-line procedure”
- “GUI procedure”

Command-line procedure

Complete the following steps to enable application support on the monitoring server for base monitoring agents, using the Linux or UNIX command-line.

Procedure

1. Start the monitoring server by running the following command:

```
./itmcmd server start tems_name
```

2. Run the following command to add the application support:

```
./itmcmd support -t tems_name pc pc pc
```

where *tems_name* is the name of the monitoring server (for example, HUB_itmserv16) and *pc* is the product code for each agent for which you want to enable application support.

By default, support seeding is skipped if application support is seeded in self-describing mode. To force support seeding over the self-describing mode seeding status, you can pass the additional parameter [-d] to the `./itmcmd support` command.

To view the product codes for the applications installed on this computer, run the following command:

```
./cinfo
```

Type 1 when prompted to display the product codes for the components installed on this computer. See your product documentation for the product code for other agents.

Activate support only for the monitoring agents for which you installed support. For example, if you installed support for the DB2 agent, you would enter the following command:

```
./itmcmd support -t tems_name ud
```

3. Stop the monitoring server by running the following command:

```
./itmcmd server stop tems_name
```

4. Restart the monitoring server by running the following command:

```
./itmcmd server start tems_name
```

GUI procedure

This section describes how to use the Manage Tivoli Enterprise Monitoring Services window on a Linux Intel or UNIX computer to enable application support on a monitoring server that is located on the local computer.

Before you begin

This procedure assumes that you have installed the support files on this computer, and that X Windows is enabled on this computer.

About this task

You can use this procedure as an alternative to the **itmcmd support** command. (This command applies only to monitoring servers that are installed on the local computer. To enable application support on a monitoring server that is located on a remote computer, see “Configuring application support on a nonlocal monitoring server from a Linux or UNIX system” on page 294.)

Complete the following steps to enable application support from the Manage Tivoli Enterprise Monitoring Services window on the local Linux or UNIX monitoring server:

Procedure

1. Log on to the computer where the Tivoli Enterprise Portal Server is installed.
2. Start the Manage Tivoli Enterprise Monitoring Services utility:
 - a. Change to the bin directory:

```
cd install_dir/bin
```
 - b. Run the following command by using the parameters described in the next table:

```
./itmcmd manage [-h ITMinstall_dir]
```

where:

-h	(optional) An option used to specify the installation directory.
<i>ITMinstall_dir</i>	The directory where the monitoring server is installed. The default installation directory is /opt/IBM/ITM.

The Manage Tivoli Enterprise Monitoring Services window is displayed.

3. Start the monitoring server if it is not already started: Right-click **Tivoli Enterprise Monitoring Server** and click **Start**.
4. Right-click **Tivoli Enterprise Monitoring Server**, and select one of the following options:
 - To enable all application support packages installed on this computer, click **Quick (all available support)**.
 - To select which application support packages you want to enable, click **Advanced . . .**
5. If you selected the **Advanced** option, the Install Product Support window is displayed. For each product-specific support package the installer checks if the Tivoli Enterprise Monitoring Server database was previously seeded with product-specific support in the self-describing mode. If so, the selected support file is excluded from the Tivoli Enterprise Monitoring Server seeding process. If you want to overwrite the support that was seeded in self-describing mode, you can select the option to **Skip self-describing mode seeding status check**.

Select the application support packages you want to install and click **Install**.

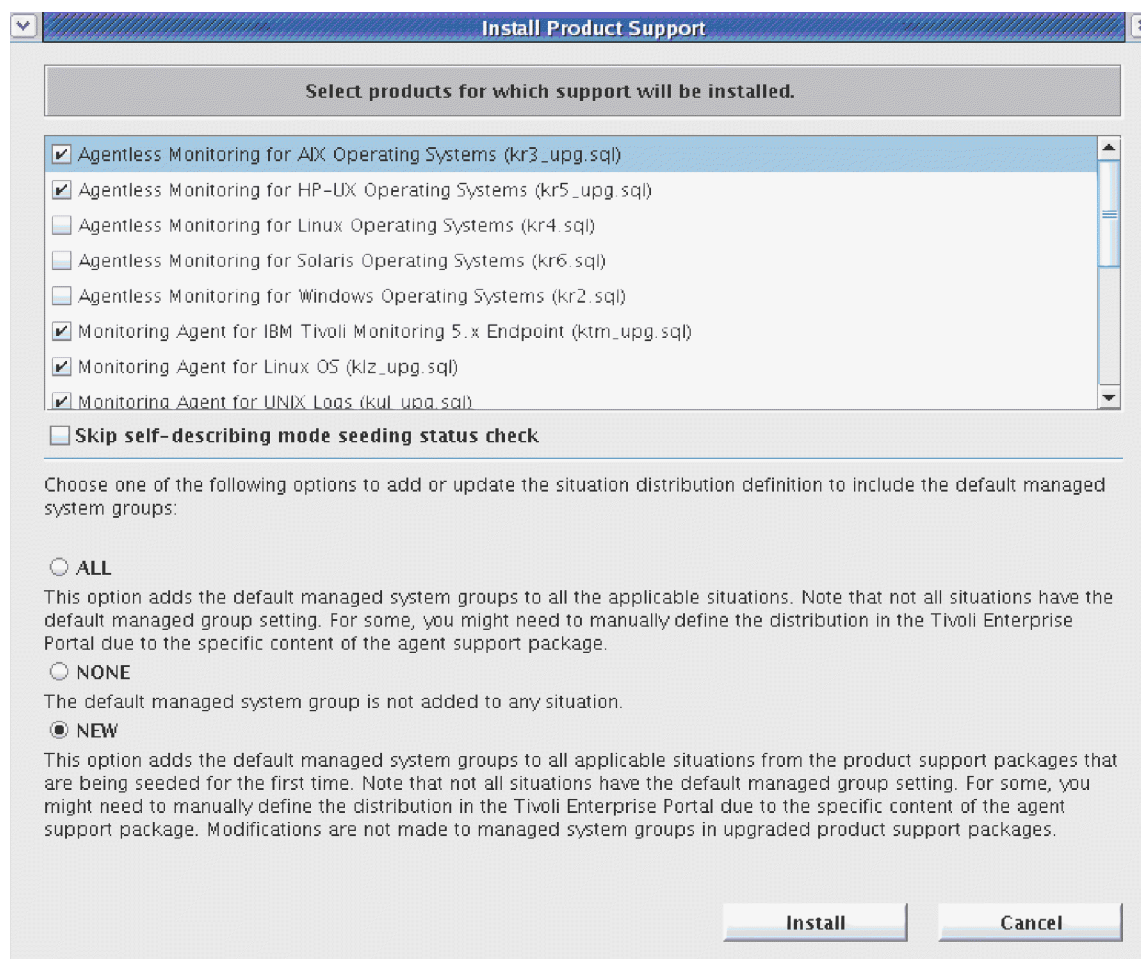


Figure 33. Install Product Support window

6. Stop and restart the monitoring server:
 - a. Right-click **Tivoli Enterprise Monitoring Server** and click **Stop**.
 - b. Right-click **Tivoli Enterprise Monitoring Server** and click **Start**.

Installing and configuring the remote monitoring servers

The following sections provide detailed information for installing and configuring a remote monitoring server:

- “Windows: Installing a remote monitoring server”
- “Linux or UNIX: Installing a remote monitoring server” on page 215

Windows: Installing a remote monitoring server

Use the following steps to install a remote monitoring server on a Windows computer.

Before you begin

IBM Tivoli Monitoring does not support multiple remote monitoring servers on the same Windows computer.

Procedure

1. Launch the installation wizard by double-clicking the setup.exe file in the IBM Tivoli Monitoring Base DVD or DVD image.

Note: If you are running Windows 2003 or Windows XP and have security set to check the software publisher of applications, you might receive an error stating that the setup.exe file is from an unknown publisher. Click **Run** to disregard this error message.

2. Click **Next** on the Welcome window.

Note: If you have another IBM Tivoli Monitoring component already installed on this computer, select **Modify** on the Welcome window to indicate that you are updating an existing installation. Click **OK** on the message telling you about preselected items. Then skip to Step 6.

3. Click **Accept** to accept the license agreement.
4. Choose the directory where you want to install the product. The default directory is C:\IBM\ITM. Click **Next**.

Note: If you specify an incorrect directory name, you will receive the following error:

The IBM Tivoli Monitoring installation directory cannot exceed 80 characters or contain non-ASCII, special or double-byte characters.

The directory name can contain only these characters:

"abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ _\.:0123456789()~-. /".

5. Type a 32-character encryption key. You can use the default key.

Note:

- a. Do not use any of the following characters in your key:

&	ampersand
	pipe
'	single quote
=	equal sign
\$	dollar sign

In addition, do not specify double-byte (DBCS) characters.

- b. Ensure that you document the value you use for the key. Use this key during the installation of any components that communicate with this monitoring server.

Click **Next** and then click **OK** to confirm the encryption key.

6. On the Select Features window, select the check box for **Tivoli Enterprise Monitoring Server**.

When you select the **Tivoli Enterprise Monitoring Server** check box, all of the check boxes in the attached subtree are automatically selected. These check boxes are for installing application support files for base monitoring agents and other to the monitoring server. (The *base* monitoring agents are included with the base IBM Tivoli Monitoring installation package.) If you leave all of the application support check boxes selected, you do not need to reconfigure application support as new agent types are added to your environment. However, installing support for many agents at a time can increase the installation time and you may still have to add support for an agent later if it has been updated. For detailed information about application support, see "Installing and enabling application support" on page 275.

Note:

- a. If you have purchased monitoring agents that run on z/OS, but have not purchased IBM Tivoli Monitoring as a separate product, expand the **Tivoli Enterprise Monitoring Server** node. Clear all check boxes in the subtree except the check boxes labeled **Tivoli Enterprise Monitoring Server** and **Summarization and Pruning Agent**.

- b. If you are updating an existing installation (you selected **Modify** on the Welcome window), all check boxes on the Select Features window reflect your choices during the initial installation. Clearing a check box has the effect of *uninstalling* the component. Clear a check box only if you want to remove a component.
7. If you want to install any agents on this remote monitoring server, expand **Tivoli Enterprise Monitoring Agents** and select the agents you want to install.
8. Click **Next** to display the Agent Deployment window.

The Agent Deployment window lists monitoring agents on this installation image that you can add to the agent depot. The agent depot contains agents that you can deploy to remote computers. For information about how to deploy agents in the agent depot to remote computers, see Chapter 11, “Deploying monitoring agents across your environment,” on page 371.

Note: By default, the agent depot is located in the *itm_installdir/CMS/depot* directory on Windows. If you want to use a different directory, create the directory (if it does not exist) and specify the directory using the DEPOTHOME key in the KBBENV file.

Select the agents, if any, that you want to add to the agent depot. (You can add agents to the agent depot at a later time by updating your installation.) Click **Next**.
9. If no IBM Tivoli Monitoring component has been previously installed on this computer, a window is displayed for you to select a program folder for the Windows Start menu. Select a program folder and click **Next**. The default program folder name is IBM Tivoli Monitoring.
10. If the TEPS Desktop and Browser Signon ID and Password window is displayed, enter and confirm the password to be used for logging on to the Tivoli Enterprise Portal. The default logon user ID, **sysadmin**, cannot be changed on this window. The logon password must match the password that you specified for **sysadmin** when you configured the hub monitoring server.

This window is not displayed if the **sysadmin** user ID has already been defined in the operating system.
11. Review the installation summary details. The summary identifies the components you are installing. Click **Next** to begin the installation.

After the components are installed, a configuration window (called the Setup Type window) is displayed.
12. Clear the check boxes for any components that are already installed and configured (at the current release level) on this computer, unless you want to modify the configuration. Click **Next** to start configuring all selected components.
13. Configure the Tivoli Enterprise Monitoring Server:
 - a. Select the type of monitoring server you are configuring: **Hub** or **Remote**. For this procedure, select **Remote**.
 - b. Verify that the name of this monitoring server is correct in the **TEMS Name** field. If it is not, change it.

The default name is *REMOTE_host_name*, for example, *REMOTE_itmserv16*. *This name must be unique in the enterprise.*
 - c. Identify the communications protocol for the monitoring server. You have four choices: IP.UDP, IP.PIPE, IP.SPIPE, or SNA. You can specify up to three methods for communication. If the method you identify as Protocol 1 fails, Protocol 2 is used as a backup. If Protocol 2 fails, Protocol 3 is used as a backup.
 - d. Leaving **TEMS Security Compatibility Mode** enabled performs the following functions:
 - 1) It allows for backwards compatibility with previous versions of IBM Tivoli Monitoring components shipped before V6.3 when securing command execution (Situation Take Action, Execute Command, Take Action, Workflow Policy Take Action). Accept the default if you have a previous version of a hub monitoring server or OMEGAMON product connecting to this remote Tivoli Enterprise Monitoring Server or if the system clocks are not synchronized across the systems with monitoring servers, the portal server, and monitoring agents.

- 2) It permits commands to still execute even if cryptographic facilities are not available. This gives you time to ensure that, for example, your z/OS Tivoli Enterprise Monitoring Server has the Integrated Cryptographic Services Facility (ICSF) available.

Note: If you disable this option and you have a mixture of pre-V6.3 and V6.3 server components, if clocks are not synchronized, or if you don't have cryptographic services available, then command executions might fail.

- e. Click **OK**.
- f. Complete the following fields for the communications protocol for the monitoring server.

Table 38. Remote monitoring server communications protocol settings

Field	Description
IP.UDP Settings: Primary Hub TEMS	
Host name or IP Address	The host name or IP address for the hub monitoring server.
Port # or Port Pools	The listening port for the hub monitoring server. The default port is 1918.
IP.PIPE Settings: Primary Hub TEMS	
Host name or IP Address	The host name or IP address for the hub monitoring server.
Port Number	The listening port for the monitoring server. The default value is 1918.
IP.SPIPE Settings: Primary Hub TEMS	
Host name or IP Address	The host name or IP address for the hub monitoring server.
Port Number	The listening port for the monitoring server. The default value is 3660.
SNA Settings: Remote TEMS	
Local LU Alias	The LU alias.
TP Name	The transaction program name for this monitoring server.
SNA Settings: Primary Hub TEMS	
Network Name	The SNA network identifier for your location.
LU Name	The LU name for the monitoring server. This LU name corresponds to the Local LU Alias in your SNA communications software.
LU 6.2 LOGMODE	The name of the LU6.2 LOGMODE. The default value is "CANCTDCS."
TP Name	The transaction program name for the monitoring server.

- g. If you are certain that you have typed the values for all of these fields with *exactly* the correct casing (upper and lower cases), you can select **Use case as typed**. However, because IBM Tivoli Monitoring is case-sensitive, consider selecting **Convert to upper case** to reduce the chance of user error.
 - h. Click **OK** to continue.
14. Enable application support on the monitoring server.
In Step 6 on page 212, you selected the base monitoring agents for which you wanted to install application support files on the monitoring server. In this step, you activate the application support through a process known as *seeding* the monitoring server.

- a. Specify the location of the monitoring server to which to add application support. You have two choices:
 - **On this computer**
 - **On a different computer**

Click **OK**.

For additional information about these parameters, press the **Help** button.

- b. Click **OK** on the Select the application support to add to the TEMS window.

This window lists the monitoring agents that you selected in Step 6 on page 212. Click **OK** to begin seeding the monitoring server (using the SQL files listed on this window).

This process can take up to 20 minutes. As the seeding process completes, a progress bar is displayed, showing the progress of seeding, in turn, the application support for the agents you selected. Once seeding completes, if support could not be added, a window is displayed showing all seeding results.

- c. Click **Next** on the message that provides results for the process of adding application support (see Figure 54 on page 284). A return code of 0 (rc: 0) indicates that the process succeeded.

Note: If the Application Support Addition Complete window is not displayed after 20 minutes, look in the IBM\ITM\CNPS\Logs\seedkpc.log files (where *pc* is the two-character product code for each monitoring agent) for diagnostic messages that help you determine the cause of the problem. For a list of product codes, see Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883.

15. Configure the communication between any IBM Tivoli Monitoring component and the monitoring server:

- a. Specify the default values for IBM Tivoli Monitoring components to use when they communicate with the monitoring server.
 - 1) If agents must cross a firewall to access the monitoring server, select **Connection must pass through firewall**.
 - 2) Identify the type of protocol that the agents use to communicate with the hub monitoring server. You have four choices: IP.UDP, IP.PIPE, IP.SPIPE, or SNA. You can specify up to three methods for communication. If the method you identify as Protocol 1 fails, Protocol 2 is used as a backup. If Protocol 2 fails, Protocol 3 is used as a backup.

Click **OK**.

- b. Complete the communication protocol fields for the monitoring server. See Table 38 on page 214 for definitions of these fields. Click **OK**.

For additional information about these parameters, press the **Help** button.

16. Click **Finish** to complete the installation.

17. Click **Finish** on the Maintenance Complete window if you are updating an existing installation.

Linux or UNIX: Installing a remote monitoring server

Use the following steps to install and configure the remote monitoring server on a Linux or UNIX computer.

Table 39. Steps for installing a remote monitoring server on a Linux or UNIX computer

Steps	Where to find information
Install the remote monitoring server. Use the same instructions as for installing the hub monitoring server.	“Installing the monitoring server” on page 201
Configure the remote monitoring server.	“Configuring the remote monitoring server” on page 216
Add application support to the remote monitoring server. Use the same instructions as for adding application support to the hub monitoring server.	“Adding application support to the hub monitoring server” on page 208

Note: Under both Linux and UNIX, IBM Tivoli Monitoring (version 6.1 fix pack 6 and subsequent) supports multiple remote monitoring servers on the same LPAR or computer (however, it does not support multiple *hub* monitoring servers on the same computer). Note that each instance of a remote monitoring server must have its own network interface card and its own unique IP address; in addition, each monitoring server must be installed on its own disk. These limitations isolate each remote monitoring server instance so you can service each one independently: you can upgrade a remote Tivoli Enterprise Monitoring Server without affecting another server's code base and shared libraries.

Configuring the remote monitoring server

Use the following steps to configure the remote monitoring server.

Procedure

1. At the command-line change to the /opt/IBM/ITM/bin directory (or the directory where you installed IBM Tivoli Monitoring).
2. Run the following command:

```
./itmcmd config -S -t tems_name
```

where *tems_name* is the name of your monitoring server (for example, remote_itmdev17).
3. Type remote to indicate that this is a remote monitoring server.
4. Press Enter to accept the default host name for the hub monitoring server. This should be the host name for hub computer. If it is not, type the correct host name and then press Enter.
5. Enter the type of protocol to use for communication with the monitoring server. You have four choices: ip.udp, ip.pipe, ip.spipe, or sna. Press Enter to use the default communications protocol (IP.PIPE).
6. If you want to set up a backup protocol, enter that protocol and press Enter. If you do not want to use backup protocol, press Enter without specifying a protocol.
7. Depending on the type of protocol you specified, provide the following information when prompted:

Table 40. UNIX monitoring server protocols and values

Protocol	Value	Definition
IP.UDP	IP Port Number	The port number for the monitoring server. The default is 1918.
SNA	Net Name	The SNA network identifier for your location.
	LU Name	The LU name for the monitoring server. This LU name corresponds to the Local LU Alias in your SNA communications software.
	Log Mode	The name of the LU6.2 LOGMODE. The default value is "CANCTDCS."
IP.PIPE	IP.PIPE Port Number	The port number for the monitoring server. The default is 1918.
IP.SPIPE	IP.SPIPE Port Number	The port number for the monitoring server. The default is 3660.

8. Press Enter to *not* specify the name of the KDC_PARTITION.

Note: You can configure the partition file at a later time, as described in Appendix C, "Firewalls," on page 861.

9. Press Enter when prompted for the path and name of the KDC_PARTITION.

10. If you want to use Configuration Auditing, press Enter. If you do not want to use this feature, type n and press Enter.
11. Press Enter to accept the default setting for the Hot Standby feature (none).
For best results, wait until after you have fully deployed your environment to configure the Hot Standby feature for your monitoring server. See the *IBM Tivoli Monitoring High Availability Guide for Distributed Systems* for information about configuring Hot Standby.
12. Press Enter to accept the default value for the Optional Primary Network Name (none).
13. Press Enter for the default **Security: Validate User** setting (NO).

Note: This option is valid only for a hub monitoring server.

14. For Tivoli Event Integration, type 2 and press Enter.

Note: This option is valid only for a hub monitoring server.

15. At the prompt asking if you want to disable Workflow Policy/Tivoli Emitter Agent event forwarding, press Enter to accept the default (2=NO).

Note: This option is valid only for a hub monitoring server.

16. To disable TEMS Security Compatibility Mode, type 2 and press Enter. Otherwise, press Enter to accept the default value (1=YES).

Leaving **TEMS Security Compatibility Mode** enabled performs the following functions:

- a. It allows for backwards compatibility with previous versions of IBM Tivoli Monitoring components shipped before V6.3 when securing command execution (Situation Take Action, Execute Command, Take Action, Workflow Policy Take Action). Accept the default if you have a previous version of a portal server, hub monitoring server, or OMEGAMON product connecting to this Tivoli Enterprise Monitoring Server or if the system clocks are not synchronized across the systems with monitoring servers, the portal server, and monitoring agents.
- b. It permits commands to still execute even if cryptographic facilities are not available. This gives you time to ensure that, for example, your z/OS Tivoli Enterprise Monitoring Server has the Integrated Cryptographic Services Facility (ICSF) available.

Note: If you disable this option and you have a mixture of pre-V6.3 and V6.3 server components, if clocks are not synchronized, or if you don't have cryptographic services available, then command executions might fail.

17. At the prompt asking if you want to configure the SOAP hubs, press Enter to save the default settings and exit the installer.

Note: This option is valid only for a hub monitoring server.

Results

The monitoring server is now configured.

A configuration file is generated in the *install_dir/config* directory with the format *host_name_ms_tems_name.config* (for example, *itmdev17_ms_HUBitmdev17.config*).

Remote copy and seeding of the application support from the hub monitoring server to distributed monitoring servers

If your requirements meet the following conditions, go to “Remote seeding from the hub monitoring server” on page 218:

- You will not be using *cinfo* to report on installed application support.
- You will not be using *itmcmd* support to perform local seeding on other monitoring servers.
- You will be using *itmcmd manage* to perform remote seeding from the hub monitoring server.

If your requirements meet the following conditions, go to “Local seeding on distributed monitoring servers” on page 219:

- You will be using `cinfo` to report on installed application support.
- You will be using `itmcmd` support to perform local seeding on other monitoring servers.
- You will not be using `itmcmd manage` to perform remote seeding from the hub monitoring server.

Remote seeding from the hub monitoring server:

You must copy the required files from the hub monitoring server to the other monitoring servers.

Before you begin

The following conditions apply for remote seeding from the hub monitoring server:

- Do not use the *FTP Catalog and Attribute files* option on the Manage Tivoli Enterprise Monitoring Services window. This option is only for sending files by FTP to z/OS systems.
- You must perform this task manually using a utility of your choice, such as FTP, SFTP, SCP, and so on.
- Replace all instances of `<TEMSNAME>` in these instructions with the name of the monitoring server on your distributed system.

Procedure

1. Copy the CAT and ATR files:
 - a. Copy the CAT files from `$CANDLEHOME/tables/ciCATrsq/RKDSCATL` on the hub monitoring server to `$CANDLEHOME/tables/ciCATrsq/RKDSCATL` on the other monitoring server system.
 - b. Copy the CAT files from `$CANDLEHOME/tables/ciCATrsq/RKDSCATL` on the other monitoring server system to `$CANDLEHOME/tables/<monitoring serverNAME>/RKDSCATL` on the same monitoring server system.
 - c. Copy the ATR files from `$CANDLEHOME/tables/ciCATrsq/ATTRLIB` on the hub monitoring server to `$CANDLEHOME/tables/ciCATrsq/ATTRLIB` on the other monitoring server system.
 - d. Copy the ATR files from `$CANDLEHOME/tables/ciCATrsq/ATTRLIB` on the other monitoring server system to `$CANDLEHOME/tables/<monitoring serverNAME>/ATTRLIB` on the same monitoring server system.
2. If you use Tivoli Enterprise Console servers, copy the .map and .baroc files:
 - a. Copy the baroc and map files from `$CANDLEHOME/tables/ciCATrsq/TECLIB` on the hub monitoring server to `$CANDLEHOME/tables/ciCATrsq/TECLIB` on the other monitoring server system.
 - b. Copy the baroc and map files from `$CANDLEHOME/tables/ciCATrsq/TECLIB` on the other monitoring server system to `$CANDLEHOME/tables/<monitoring serverNAME>/TECLIB` on the same monitoring server system.
3. Use the Manage Tivoli Enterprise Monitoring Services window to perform the seeding. On Linux/UNIX systems, run the `itmcmd manage` command to display the Manage Tivoli Enterprise Monitoring Services window:
 - a. In the Manage Tivoli Enterprise Monitoring Services window, select the Tivoli Enterprise Monitoring Server.
 - b. Click **Actions** → **Install Product Support**.
 - c. A pop-up window is displayed with the options **On this computer** and **On a different computer**. Click **On a different computer** and then click **OK**.
 - d. Click **OK** on the pop-up window after you have ensured that the monitoring server on the other machine is running.
 - e. The **Non-Resident monitoring server Connection** window is displayed. Complete the information in this window and click **OK**.
 - f. A second **Non-Resident monitoring server Connection** window is displayed. Enter the information required and click **OK**.

- g. If all of the connection information is correct a list of products available for seeding is displayed. Select the desired products from the list and execute the seeding.
- h. Remember to recycle the monitoring server on the system that has been updated.

Local seeding on distributed monitoring servers:

You must copy the required files from the hub monitoring server to the other monitoring servers.

Before you begin

The following conditions apply for local seeding on distributed monitoring servers:

- Do not use the *FTP Catalog and Attribute files* option on the Manage Tivoli Enterprise Monitoring Services window. This option is only for sending files by FTP to z/OS systems.
- You must perform this task manually using a utility of your choice, such as FTP, SFTP, SCP, and so on.
- Replace all instances of <monitoring serverNAME> in these instructions with the name of the monitoring server on your distributed system.

Procedure

1. Copy the CAT and ATR files:
 - a. Copy the CAT files from \$CANDLEHOME/tables/ciCATrsq/RKDSCATL on the hub monitoring server to \$CANDLEHOME/tables/ciCATrsq/RKDSCATL on the other monitoring server system.
 - b. Copy the ATR files from \$CANDLEHOME/tables/ciCATrsq/ATTRLIB on the hub monitoring server to \$CANDLEHOME/tables/ciCATrsq/ATTRLIB on the other monitoring server system.
2. If you use Tivoli Enterprise Console servers, copy the .map and .baroc files:
 - Copy the baroc and map files from \$CANDLEHOME/tables/ciCATrsq/TECLIB on the hub monitoring server to \$CANDLEHOME/tables/ciCATrsq/TECLIB on the other monitoring server system.
3. Copy the *tms.ver files to be able to use cinfo and itmcmd support.
 - Copy the *tms.ver files from \$CANDLEHOME/registry on the hub monitoring server to \$CANDLEHOME/registry on the other monitoring server system.
4. Perform local seeding on the other monitoring server system:
 - a. Log in to the system on which the other monitoring server resides.
 - b. Change directory to the \$CANDLEHOME/bin directory.
 - c. Run the following command:


```
./itmcmd support -t <monitoring serverNAME> pc
```

where *pc* can be a space delimited list to seed for multiple products in one run.

- d. Remember to recycle the monitoring server on the system that has been updated.

Installing and configuring Tivoli Enterprise Monitoring Automation Server

You must install and configure the Tivoli Enterprise Monitoring Automation Server if you are using the Registry Services component of IBM Jazz for Service Management and other products to integrate with IBM Tivoli Monitoring using OSLC linked data interfaces. Before installation, you must be aware of the existing Registry Services and Performance Monitoring service provider configuration details.

For more information on setting up Registry Services and optionally Security Services before installing the monitoring automation server, see “Required software and memory requirements for using the OSLC Performance Monitoring service provider” on page 94.

The Performance Monitoring server provider does not support the dynamic refresh of application support files. This means that the Tivoli Enterprise Monitoring Automation Server must be recycled when new application support is added to a system, regardless of whether the support was added through a self-describing agent, a regular installation, or a manual copy. However, if the monitoring agent does not provide OSLC support, then the Automation Server does not need to be recycled. Typically, agents before IBM Tivoli Monitoring V6.3 do not provide OSLC support. See the agent documentation to determine whether an agent provides OSLC support.

Note: The Tivoli Enterprise Monitoring Automation Server cannot be used if the Hot Standby option is enabled on the Hub monitoring server. The Automation Server will terminate if it connects to a Hub monitoring server that is configured in Hot Standby (FTO) mode.

The following sections provide detailed information for installing and configuring a Tivoli Enterprise Monitoring Automation Server:

- “Windows: Installing a Tivoli Enterprise Monitoring Automation Server”
- “Linux or UNIX: Installing a Tivoli Enterprise Monitoring Automation Server” on page 222

Windows: Installing a Tivoli Enterprise Monitoring Automation Server

Use the following steps to install the automation server on a Windows computer.

Before you begin

A locally installed hub monitoring server is a prerequisite for the Tivoli Enterprise Monitoring Automation Server. An error message reports a discrepancy if you attempt to configure the automation server with a remote monitoring server. For a pristine installation, when you select **Tivoli Enterprise Monitoring Automation Server** on the Select Features window, the **Tivoli Enterprise Monitoring Server** check box is automatically selected. You must configure the automation server connection information before you can start it.

Procedure

1. Launch the installation wizard by double-clicking the setup.exe file in the IBM Tivoli Monitoring Base DVD or DVD image.
2. Click **Next** on the Welcome window.
3. Select the **I accept the terms in the license agreement** radio button and click **Next**.
4. Specify the directory where you want to install the software and accompanying files. The default location is C:\IBM\ITM. Click **Next**.
5. Type an encryption key to use. The key required is the same as that used for the installation of the hub monitoring server. Click **Next** and then **OK** to confirm the encryption key.
6. On the Select Features window, select **Tivoli Enterprise Monitoring Automation Server** from the list of components to install.

If a Tivoli Enterprise Monitoring Server is not already installed the **Tivoli Enterprise Monitoring Server** is selected by default when you select **Tivoli Enterprise Monitoring Automation Server**. The check boxes in the subtree under Tivoli Enterprise Monitoring Server for installing application support files for monitoring agents are not selected by default. Select the monitoring agent application support files that you want to install to the monitoring server. For detailed information about application support, see “Installing and enabling application support” on page 275.

7. Click **Next** to display the Agent Deployment window.
8. Select the agents, if any, that you want to add to the agent depot. (You can add agents to the agent depot at a later time by updating your installation.) Click **Next**.

If you already have a **Tivoli Enterprise Monitoring Server** installed, proceed to step 11 on page 221.

9. If no **Tivoli Enterprise Monitoring Server** is installed on this computer, a window is displayed for you to select a program folder. Click **Next** to accept the default program folder name. If you want to use a different program folder name, type the program folder name and click **Next**.
10. In the TEPS Desktop and Browser Signon ID and Password window, specify the user ID and password to be used for logging on to the Tivoli Enterprise Portal and click **Next**.
11. Review the installation summary details. The summary identifies the components you are installing. Click **Next** to begin the installation.

After the components are installed, a configuration window (called the Setup Type window) is displayed.

12. Select the check box for **Configure Tivoli Enterprise Monitoring Server**. Clear the check boxes for any components that are already installed and configured (at the current release level) on this computer, unless you want to modify the configuration. Click **Next** to display the Hub TEMS Configuration window.
13. Select **Hub** as the type of monitoring server you are configuring and click **OK** to continue. Press the **Help** button for additional information about the configuration parameters of the monitoring server.
14. You are asked if you want to configure the automation server. Click **Yes** to display the **Configure Tivoli Enterprise Monitoring Automation Server** window.
15. Complete your Registry Services configuration details and automation server connection details. When you select **Configure Registry Services connection**, you must specify the correct Protocol, Hostname, and Port number for the WebSphere server where Registry Services is installed.

The following criteria apply to the **User** and **Password** fields:

- Providing both a user ID and a password is accepted.
- Leaving both the user ID and a password fields blank is accepted.
- Providing one parameter without the other is not accepted.

In the **Automation Server connection information** section, provide the HTTP protocol and Port number for clients to connect to the automation server. Click **OK** to continue.

Figure 34. Configure Tivoli Enterprise Monitoring Automation Server

16. Click **OK** on the Hub TEMS Configuration window.
17. Click **Finish** to complete the installation.

18. Click **Finish** on the Maintenance Complete window if you are updating an existing installation.

Results

The Manage Tivoli Enterprise Monitoring Services utility is opened. (This might take a few minutes.)

During pristine installation of the automation server, after it is successfully configured, it is automatically started along with other IBM Tivoli Monitoring components. If the hub monitoring server is unconfigured, or is switched from hub to remote, the automation server is automatically stopped and unconfigured.

What to do next

When the automation server is started, the Performance Monitoring service provider registers itself with Registry Services and starts registering resources for managed systems when they come online and the agent has an OSLC agent template.

To configure the Performance Monitoring service provider to use the Security Services component of Jazz for Service Management to authenticate OSLC client applications, you must set the `KAS_SECURITY_SERVICES_ENABLED` environment variable to Yes in the Tivoli Enterprise Monitoring Automation Server's environment file and restart the monitoring automation server.

After you install application support for agents that have an OSLC template, restart the automation server. You must restart the automation server if the application support is manually installed or is installed using the self-describing agent feature. After performing these steps, IBM Tivoli Monitoring is ready to integrate with other products using OSLC linked data interfaces.

Linux or UNIX: Installing a Tivoli Enterprise Monitoring Automation Server

Use the following steps to install the automation server on a Linux or UNIX computer.

Before you begin

A Tivoli Enterprise Monitoring Automation Server requires a locally installed hub monitoring server (see “Installing and configuring the hub Tivoli Enterprise Monitoring Server” on page 196).

Procedure

1. In the directory where you extracted the installation files, run the following command:
`./install.sh`
2. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default (/opt/IBM/ITM). If you want to use a different installation directory, type the full path to that directory and press Enter.

Note: You must not specify the path of the directory containing `./install.sh` as your IBM Tivoli Monitoring home directory. On certain platforms, this can cause the plug-in JAR files to overwrite themselves and become zero length files. The installation will fail as a result.

In addition, do not use the number sign (#) or exclamation mark (!) characters anywhere in your installation directory path.

3. If the directory you specified does not exist, you are asked whether to create it. Type 1 to create this directory.
4. The following prompt is displayed:

Select one of the following:

- 1) Install products to the local host.
- 2) Install products to depot for remote deployment (requires TEMS).
- 3) Install TEMS support for remote seeding
- 4) Exit install.

Please enter a valid number:

Type **1** to start the installation and press Enter.

The end user license agreement is displayed. Press Enter to read through the agreement.

5. Type **1** to accept the agreement and press Enter.

If you are upgrading IBM Tivoli Monitoring from global GSKit enabled to local GSKit enabled, the following text is displayed:

Local GSKit has been installed for IBM Tivoli Monitoring. The global GSKit installation will no longer be used by IBM Tivoli Monitoring. It will not be uninstalled automatically, because it might be used by other applications. You may uninstall it manually, if IBM Tivoli Monitoring is the only application utilizing it.

6. Type a 32-character encryption key and press Enter. If you want to use the default key, press Enter without typing any characters.

Note:

- a. Do not use any of the following characters in your key:

&	Ampersand
	Pipe
'	Single quote
=	Equal sign
\$	Dollar sign

In addition, do not specify double-byte (DBCS) characters.

- b. Ensure that you document the value you use for the key. Use this key during the installation of any components that communicate with this monitoring server.

Note: If you are upgrading IBM Tivoli Monitoring instead of a pristine installation, you are prompted by the installer to upgrade the common prerequisite components for agents and servers. The product packages available for this operating system and component support categories are listed.

7. Type the number that corresponds to your operating environment. For a new installation, type **1** IBM Tivoli Monitoring components for this operating system. This option installs all components for your operating system.
8. Type **1** to confirm the selection of product packages in step 7.
9. Type the number that corresponds to the **Tivoli Enterprise Monitoring Automation Server**. Press Enter.

Note: The automation server requires a locally installed hub monitoring server. If no monitoring server is installed, type the number that corresponds to the Tivoli Enterprise Monitoring Server.

10. Type **1** to confirm the installation.

Attention: If you are installing the Tivoli Enterprise Monitoring Server at this time, the monitoring server application support is also installed. You are prompted to answer the following questions as the installation completes:

- Please enter TEMS name [TEMS is default]:

- Do you want to install additional products or product support packages?
[1=Yes, 2=No; default is "2"]

- Do you want to seed product support on the Tivoli Enterprise Monitoring Server?
[1=Yes, 2=No; default is "1"]

- Type your selection here [Number "3" or "NEW" is default]:

11. If your IBM Tivoli Monitoring environment is not already secured you will be asked at this point if you want to secure it. [1=Yes, 2=No; default is "1"]. If your IBM Tivoli Monitoring environment is already secured this question is skipped. The product installation process creates the majority of directories and files with world write permissions. IBM Tivoli Monitoring provides the secureMain utility to help you keep the monitoring environment secure. You can secure your installation now, or manually execute the secureMain utility later. For more information, see Appendix G, "Securing your IBM Tivoli Monitoring installation on Linux or UNIX," on page 925

What to do next

The next step is to configure your automation server.

After you have configured your automation server, you can perform the following steps:

- To configure the Performance Monitoring service provider to use the Security Services component of Jazz for Service Management to authenticate OSLC client applications, you also must set the KAS_SECURITY_SERVICES_ENABLED environment variable to Yes in the Tivoli Enterprise Monitoring Automation Server's environment file and restart the monitoring automation server.
- After you install application support for agents that have an OSLC template, restart the automation server. You must restart the automation server if the application support is manually installed or is installed using the self-describing agent feature.

After performing these steps, IBM Tivoli Monitoring is ready to integrate with other products using OSLC linked data interfaces.

Configuring the Tivoli Enterprise Monitoring Automation Server

Use one of the following procedures to configure the Tivoli Enterprise Monitoring Automation Server.

- "Command-line procedure"
- "GUI procedure" on page 225

Command-line procedure

Complete the following steps to configure the Tivoli Enterprise Monitoring Automation Server, by using the Linux or UNIX command-line.

Procedure

1. At the command-line, change to the /opt/IBM/ITM/bin directory (or the directory where you installed IBM Tivoli Monitoring).
2. Run the following command:

```
./itmcmd config -A as
```

where **as** is the product code of the Tivoli Enterprise Monitoring Automation Server.

3. Enter your Registry Services connection settings. When you set the Edit Registry Services connection settings parameter to Yes, you must specify the correct Protocol, Hostname, and Port

number for the WebSphere server where Registry Services is installed. When you set the Edit Registry Services connection user settings parameter to Yes, you must enter a user ID and Password for the Registry Services

Note: You must select HTTPS as the protocol if you are using Security Services in addition to Registry Services.

This section contains an example of the command-line syntax. In this example, you must replace *server* with the host name for your Registry Services provider registry:

```
Set Registry Services connection settings. Selecting No will clear any previously set
Registry Services connection settings. [1=Yes, 2=No](Default is: 1):
Protocol [1=HTTP, 2=HTTPS](Default is: 1):
Hostname(Default is: ): server
HTTP Port Number(Default is: 16310):
Set Registry Services user and password. Selecting No will clear any previously set
Registry Services user and password settings. [1=Yes, 2=No](Default is: 1):
User(Default is: ): user
Password:
Re-type: Password:
```

- 4. Provide the HTTP protocol and port for clients to connect to the automation server.
Tivoli Enterprise Monitoring Automation Server connection information
Protocol [1=HTTP, 2=HTTPS](Default is: 1):
HTTP Port Number(Default is: 10001):

Results

The automation server is now configured.

What to do next

Manually start the automation server:

```
./itmcmd agent start as
```

You can stop the automation server with the following command:

```
./itmcmd agent stop as
```

GUI procedure

Complete these steps to use the Manage Tivoli Enterprise Monitoring Services window on a Linux Intel or UNIX computer to configure the Tivoli Enterprise Monitoring Automation Server.

Procedure

- 1. Log on to the computer where the automation server is installed.
- 2. Start the Manage Tivoli Enterprise Monitoring Services utility:
 - a. Change to the bin directory:
`cd install_dir/bin`
 - b. Run the following command by using the parameters described in the next table:
`./itmcmd manage [-h ITMinstall_dir]`

where:

-h	(optional) An option used to specify the installation directory.
ITMinstall_dir	The directory where the monitoring server is installed. The default installation directory is /opt/IBM/ITM.

The Manage Tivoli Enterprise Monitoring Services window is displayed.

3. Right-click **Tivoli Enterprise Monitoring Automation Server** and click **Configure**.
4. Complete your Registry Services connection details and automation server connection details. When you select **Configure Registry Services connection**, you must specify the correct Protocol, Hostname, and Port number for the WebSphere server where Registry Services is installed.

Note: You must select HTTPS as the protocol if you are using Security Services in addition to Registry Services.

The following criteria apply to the **User** and **Password** fields:

- Providing both a user ID and a password is accepted.
- Leaving both the user ID and a password fields blank is accepted.
- Providing one parameter without the other is not accepted.

In the **Automation Server connection information** section, provide the HTTP protocol and port for clients to connect to the automation server. Click **OK** to continue.

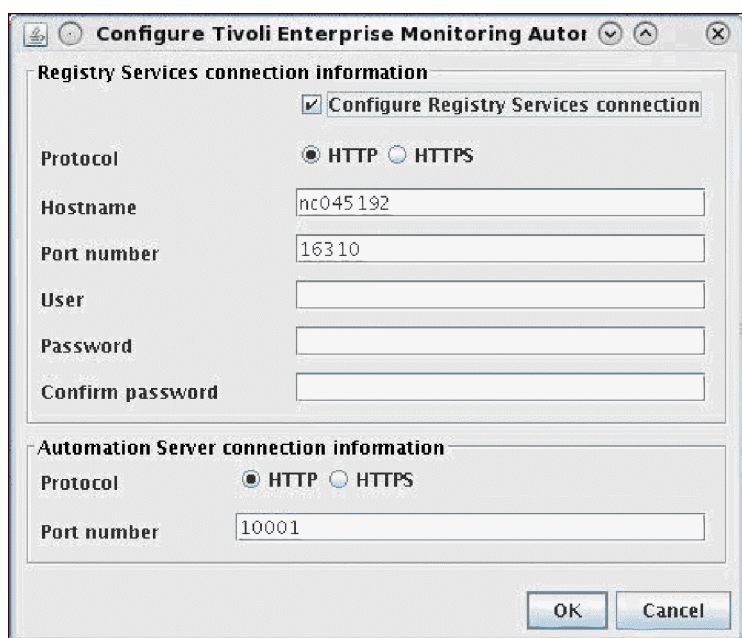


Figure 35. Configure Tivoli Enterprise Monitoring Automation Server

5. If you reconfigure the automation server when it is already running you are asked if you want to restart the automation server or not.
You can also stop and restart the automation server from the Manage Tivoli Enterprise Monitoring Services window:
 - a. Right-click **Tivoli Enterprise Monitoring Automation Server** and click **Stop**.
 - b. Right-click **Tivoli Enterprise Monitoring Automation Server** and click **Start**.

Installing the Tivoli Enterprise Portal Server

You can install the Tivoli Enterprise Portal Server on Windows, Linux or AIX.

Follow the instructions for your operating system:

- “Windows: Installing the portal server” on page 227
- “Linux or AIX: Installing the portal server” on page 241

Note:

1. In IBM Tivoli Monitoring V6.3, IBM Installation Manager installs the following sub components during a Tivoli Enterprise Portal Server installation:
 - IBM HTTP Server V8.0 for Tivoli Enterprise Portal Server
 - Web Server Plug-ins for IBM WebSphere Application Server V8.0 for Tivoli Enterprise Portal Server

These sub components are required by the portal server and must not be uninstalled manually by using the IBM Installation Manager. Otherwise, the portal server can not operate. The sub components will be automatically uninstalled if you uninstall the portal server.

2. IBM Installation Manager version 1.6.1 is installed (if it is not already installed on your machine) at the same time as a version 6.3 installation or upgrade of the Tivoli Enterprise Portal Server. For information about the default directories where Installation Manager is installed and the directory that is used to store information about installed packages, see Default installation locations for application files and Agent data location in the IBM Installation Manager information center. If you want to install Installation Manager in a different location, download Installation Manager version 1.6.1 (or later) and manually install it in a directory of your choice before installing or upgrading the Tivoli Enterprise Portal Server. However, IBM Tivoli Monitoring does not support an agent data location other than the default directory used by IBM Installation Manager. To download IBM Installation Manager, go to <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.
3. Only a single instance of IBM Installation Manager can run at any one time. If IBM Installation Manager is already installed on the server machine, make sure that it is not running before you begin a version 6.3 installation or upgrade of the Tivoli Enterprise Portal Server.
4. An installation or upgrade might fail for a 32-bit Tivoli Enterprise Portal Server on a 64-bit Linux operating system that contains version 1.6.0 of 64-bit IBM Installation Manager. The failure occurs because the installation or upgrade of the Tivoli Enterprise Portal Server V6.3 triggers an installation or upgrade of the IBM Installation Manager to version 1.6.1. The 32-bit IBM Tivoli Monitoring V6.3 installation media only contains a 32-bit version of IBM Installation Manager, and therefore cannot proceed with the upgrade of the existing 64-bit IBM Installation Manager. To resolve this problem you must manually update the 64-bit IBM Installation Manager to version 1.6.1 (or later) before beginning the installation or upgrade of Tivoli Enterprise Portal Server V6.3. To download IBM Installation Manager, go to <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.
5. When you install or upgrade the Tivoli Enterprise Portal Server, the Tivoli Enterprise Services User Interface Extensions software is automatically installed in the same directory. The portal server extensions are required for some products that use the Tivoli Enterprise Portal (for example, IBM Tivoli Composite Application Manager for Service Oriented Architecture).

The Tivoli Enterprise Services User Interface Extensions software is supported on the same operating systems as the Tivoli Enterprise Portal Server.

Windows: Installing the portal server

Complete the steps in this section to install the Tivoli Enterprise Portal Server and portal client on a Windows computer.

About this task

The installation procedure for a portal server on Windows includes steps for configuring the connection between the portal server and the following components:

- The hub monitoring server
- The portal server database
- The Tivoli Data Warehouse database
- Tivoli Authorization Policy Server

Note:

1. If you have not set up the Tivoli Data Warehouse, complete this procedure but click **No** when asked if you want to configure the connection to the data warehouse. You can reconfigure the connection after you set up the warehouse. See Step 20 on page 235 for more information.
2. The Windows userid you use when creating the portal server database must be a **local** ID with Administrator privileges. In other words, it cannot be a domain ID.

Procedure

1. Launch the installation wizard by double-clicking the setup.exe file in the IBM Tivoli Monitoring Base DVD or DVD image.
2. Click **Next** on the Welcome window.
3. Read and accept the software license agreement by clicking **Accept**.
4. Specify the directory where you want to install the portal server software and accompanying files. The default location is C:\IBM\ITM. Click **Next**

Note: If you specify an incorrect directory name, you will receive the following error:

The IBM Tivoli Monitoring installation directory cannot exceed 80 characters or contain non-ASCII, special or double-byte characters.

The directory name can contain only these characters:

"abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ _\.:0123456789()~-. /".

5. Type an encryption key to use. This key should be the same as what was used during the installation of the hub monitoring server to which this portal server will connect. Click **Next** and then **OK** to confirm the encryption key.
6. On the Select Features window, select **Tivoli Enterprise Portal Server** from the list of components to install.

When you select the **Tivoli Enterprise Portal Server** check box, all of the check boxes in the attached subtree are automatically selected. The support check boxes in the subtree are for installing application support files for base monitoring agents to the monitoring server. (The *base* monitoring agents are included with the base IBM Tivoli Monitoring installation package.) It is best to leave all of the support check boxes selected so you do not need to reconfigure application support as new agent types are added to your environment. For detailed information about application support, see "Installing and enabling application support" on page 275.

Note:

- a. If you have purchased monitoring agents that run on z/OS, but have not purchased IBM Tivoli Monitoring as a separate product, expand the **Tivoli Enterprise Portal Server** node. Clear all check boxes in the subtree except the check boxes labeled **Tivoli Enterprise Portal Server** and, optionally, **TEC GUI Integration**. (See Step 7.a.)
 - b. If you are updating an existing installation, all check boxes on the Select Features window reflect your choices during the initial installation. If your installation media contains a newer version of a previously installed component, the component will be upgraded unless you clear the check box for that component.
 - c. The Eclipse Help Server is automatically selected when you select the Tivoli Enterprise Portal Server.
7. Optionally select the following additional components to install:
 - a. If you want to view events on the IBM Tivoli Enterprise Console event server through the Tivoli Enterprise Portal, expand **Tivoli Enterprise Portal Server** and ensure that **TEC GUI Integration** is selected.
 - b. If you want to install a portal desktop client on this computer, select **Tivoli Enterprise Portal Desktop Client**.

When you select the **Tivoli Enterprise Portal Desktop Client** check box, all of the check boxes in the attached subtree are automatically selected. These check boxes are for installing application

support files for base monitoring agents to the portal desktop client. Leave these check boxes selected as you did for the portal server in Step 6 on page 228.

Note: If you have purchased monitoring agents that run on z/OS, but have not purchased IBM Tivoli Monitoring as a separate product, expand the **Tivoli Enterprise Portal Desktop Client** node. Clear all check boxes in the subtree except the check boxes labeled **Tivoli Enterprise Portal Desktop Client** and, optionally, **TEC GUI Integration**.

8. Click **Next**.

9. If a monitoring server is not installed on this computer, go to Step 10.

If you are installing the portal server on a computer that already has a monitoring server installed, the Agent Deployment window is displayed.

The Agent Deployment window lists monitoring agents on this installation image that you can add to the agent depot. The agent depot contains agents that you can deploy to remote computers. For information about how to deploy agents in the agent depot to remote computers, see Chapter 11, “Deploying monitoring agents across your environment,” on page 371.

Note: By default, the agent depot is located in the *itm_installdir/CMS/depot* directory on Windows. If you want to use a different directory, create the directory (if it does not exist) and specify the directory using the DEPOTHOME key in the KBBENV file.

Select the agents, if any, that you want to add to the agent depot. (You can add agents to the agent depot at a later time by updating your installation.) Click **Next**.

10. If no IBM Tivoli Monitoring component has been previously installed on this computer, a window is displayed for you to select a program folder for the Windows Start menu. Select a program folder and click **Next**. The default program folder name is IBM Tivoli Monitoring.
11. If you are installing the Tivoli Enterprise Monitoring Server at the same time as the Tivoli Enterprise Portal Server, you are asked to set a sign on password for the Tivoli Enterprise Portal desktop and browser clients.
12. Review the installation summary details. The summary identifies what you are installing and where you chose to install it. Click **Next** to start the installation.

Important: Only a single instance of IBM Installation Manager can run at any one time. A message is displayed if the installer detects another running instance of IBM Installation Manager. You are prompted to close IBM Installation Manager and then click **OK** to continue with the installation. When you click **OK** the installer rechecks that IBM Installation Manager is not already running. You can click **Cancel** to exit the installation.

After installation is complete, a configuration window (called the Setup Type window) is displayed.

13. Clear the check boxes for any components that are already installed and configured (at the current release level) on this computer, unless you want to modify the configuration. (For example, clear the check box for the Tivoli Enterprise Monitoring Server if it has already been installed and configured on this computer.) Click **Next** to start configuring all selected components.
14. You are asked to choose the database management system you want to use for your portal server database, Derby, DB2 Database for Linux, UNIX, and Windows, or Microsoft SQL Server, as shown in Figure 36 on page 230. Note that if a particular Database Type is uninstalled on this computer or if it installed but not at the necessary release level, the radio button is grayed out.

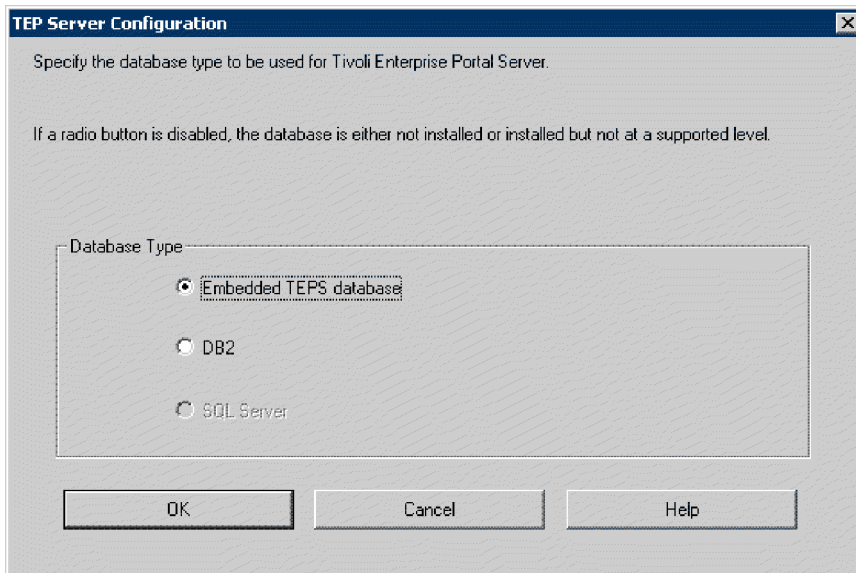


Figure 36. The Select Database for Tivoli Enterprise Portal window

15. Select the communications protocol for the connection between the portal server and the hub monitoring server. You can choose from the following protocol options:

- IP.PIPE
- IP.SPIPE
- SNA
- IP.UDP

You can select up to three protocols and specify if the connection between the portal server and the hub monitoring server passes through a firewall. Click **OK**.

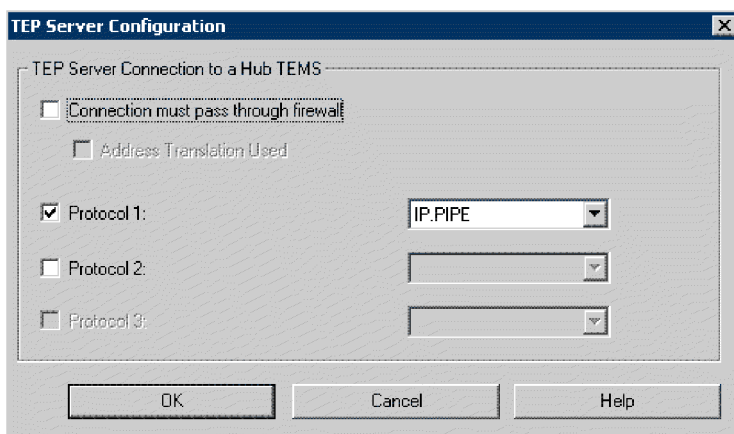


Figure 37. Communications protocol between the portal server and the hub monitoring server

16. Configure communications for the Tivoli Enterprise Portal Server:
 - a. Type the host name of the computer where you are installing the portal server. (The host name of this computer is displayed by default.)
 - b. Select **Validate user with LDAP** if you want the portal server to use an LDAP registry to authenticate users and, optionally, to enable single signon. A best practice is to install and configure the portal server with LDAP validation disabled. After a basic installation of IBM Tivoli

Monitoring has been completed and tested, you can reconfigure the portal server and enable LDAP user validation using the instructions in the *Enabling user authentication* chapter of the *IBM Tivoli Monitoring Administrator's Guide*.

Note: If you enable LDAP user validation, the LDAP server must be started so that Tivoli Monitoring can verify a connection to the LDAP server. If a connection cannot be established, LDAP user validation will be disabled when the portal server installation completes.

- c. Select **Enable Single Sign On** if you enabled LDAP user validation and you want the portal server to use single signon with other applications such as IBM Dashboard Application Services Hub or with other applications that launch the portal client. A best practice is to install and configure the portal server with single signon disabled. After a basic installation of IBM Tivoli Monitoring has been completed and tested, you can reconfigure the portal server and enable single signon using the instructions in the *Enabling user authentication* chapter of the *IBM Tivoli Monitoring Administrator's Guide*.
- d. Select the **Enable the dashboard data provider** check box if Dashboard Application Services Hub will be used with monitoring dashboard applications such as IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM, or will be used to create or enable custom dashboard pages for monitoring agents. (See the *Preparing your dashboard environment* topic in the *IBM Tivoli Monitoring Administrator's Guide* for details on additional configuration tasks to perform after installation.)
- e. The domain override value is optional and is only valid when the dashboard data provider is enabled. It changes the default dashboard data provider ID and domain name for authorization policies from itm.<Hub monitoring server name> to itm.<domain override value>. The value may not exceed 124 characters. You should configure a domain override value for these scenarios:
 - The Hot Standby high availability feature is being used for the Hub monitoring server. By configuring a domain override value, the dashboard data provider ID and domain name will not change when the portal server is configured to connect to the new acting Hub monitoring server. If you do not configure a domain override value in this scenario, you must reconfigure the connection between the IBM Dashboard Application Services Hub and the dashboard data provider and update any domain specific authorization policies when the portal server is configured to connect to the new acting Hub monitoring server.
 - You have multiple Hub monitoring servers that are using a common set of authorization policies for controlling dashboard access and you want to create some domain specific authorization policies. You should specify a domain override value for this scenario if you want to use a more user friendly domain name in your authorization policies than the default value of itm.<Hub monitoring server name>.
 - If you modify the domain override after you have configured a connection in your Dashboard Application Services Hub Server to the dashboard data provider then you must delete the connection and re-add it. See the *IBM Tivoli Monitoring Administrator's Guide* for details on how to configure a dashboard data provider connection. Also if you have created any domain specific authorization policies using the default domain name then you must delete the permissions that use the previous domain name and create new permissions that use the new domain name when you change the domain override value. See the *IBM Tivoli Monitoring Command Reference* for details on the tivcmd CLI commands used to create and work with authorization policies.
- f. Select the **Enable authorization policies** check box if you want to use authorization policies to control which managed systems and managed system groups a user can access in monitoring dashboards of the IBM Dashboard Application Services Hub. Once authorization policies are enabled, only authorized users will see data in monitoring dashboards. You should only enable authorization policies if your administrators have already created the initial set of policies for dashboard user access using the Tivoli Authorization Policy Server and tivmd Command Line

Interface for Authorization Policy. See the *Using role-based authorization policies* topic in the *IBM Tivoli Monitoring Administrator's Guide* for details on how to create and work with authorization policies.

Note: If authorization policies have not been created yet, leave this box unchecked. After the authorization policies are ready reconfigure the portal server and enable authorization policies.

g. Click **OK**.

TEP Server Configuration

IP/UDP Settings of the TEMS

Hostname or IP Address: WIN-7UUCT1441A5

Port number and/or Port Pools: 1918

IP/PIPE Settings of the TEMS

Hostname or IP Address: WIN-7UUCT1441A5

Port number: 1918

IP/SPIPE Settings of the TEMS

Hostname or IP Address: WIN-7UUCT1441A5

Port number: 3660

SNA Settings of the TEMS

Network Name:

LU Name:

LU6.2 LOGMODE: CANCTDCS

TP Name: SNASOCKETS

Local LU Alias:

(LU Alias is not required if using default)

Entry Options

☐ Use case as typed

☒ Convert to upper case

LDAP Security

☐ Validate User with LDAP ?

☐ Enable Single Sign On

Dashboard data provider

☒ Enable the dashboard data provider

Domain override:

☒ Enable authorization policies

NAT Settings OK Cancel Help

Figure 38. Tivoli Enterprise Portal Server Configuration

17. If you selected **Validate User with LDAP** in step 16.b, enter the LDAP configuration details in the fields provided and click **OK**. For more information about these parameters, see *Prerequisites for configuring LDAP authentication on the portal server* in the *IBM Tivoli Monitoring Administrator's Guide*.

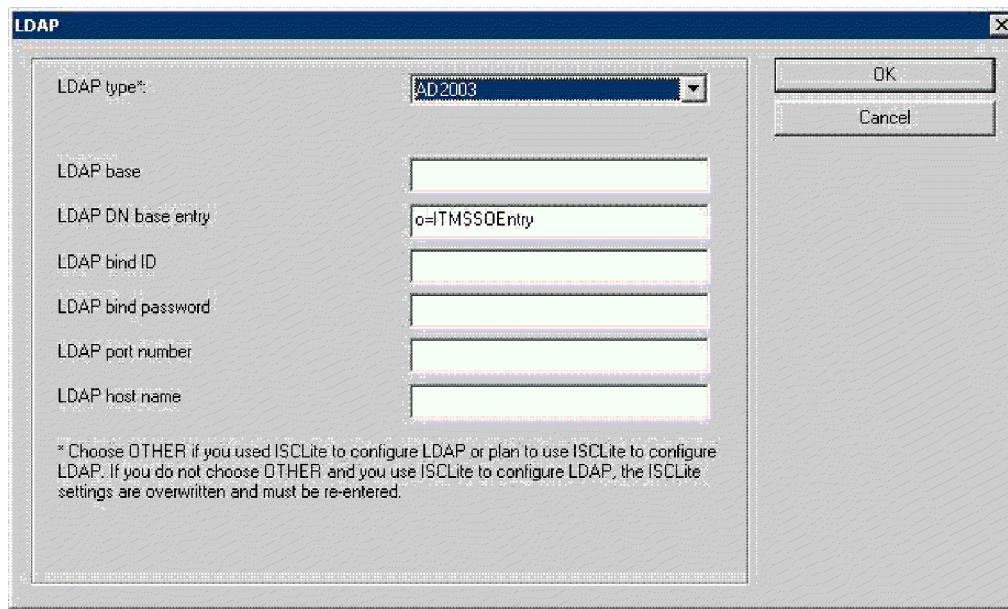


Figure 39. LDAP configuration

18. If you selected **Enable Single Sign On** in step 16.c the Single Sign On window is displayed. Enter values for the following fields and click **OK**:

- Domain name
- Realm name

You can also export or import LTPA keys from other participating single sign-on applications. For more information about configuring these fields and setting up SSO, see *Prerequisites for configuring LDAP authentication on the portal server* and *About single sign-on* in the *IBM Tivoli Monitoring Administrator's Guide*.

19. If you selected **Enable authorization policies** in step 16.f, the Authorization Policy Server Configuration window is displayed. To avoid time-consuming authorization checks across the network, the Tivoli Enterprise Portal Server retrieves its own local copy of the Authorization Policy Server's data store. For the data store retrieval to succeed, the Tivoli Enterprise Portal Server requires several pieces of information about the Authorization Policy Server.

Authorization Policy Server Configuration

Authorization Policy Server connection information

Hostname or IP address
myPSserver.ibm.com

Protocol
☒ HTTPS
☐ HTTP

Port
16311

Polling interval (minutes)
30

Policy store expiration interval
7 days 0 hours

User ID
myPSuser

Password

Confirm password

OK Cancel Help

Figure 40. Authorization Policy Server Configuration

The fields on the configuration window are described in the following table:

Table 41. Configuration information for the Authorization Policy Server

Field	Description
Hostname or IP Address	IP Address or fully qualified hostname of the IBM Dashboard Application Services Hub with the Authorization Policy Server. This parameter is required.
Protocol	Choose the protocol used to connect to the IBM Dashboard Application Services Hub with the Authorization Policy Server. The default value is HTTPS. This parameter is not required.
Port	Choose the port used to connect to the IBM Dashboard Application Services Hub with the Authorization Policy Server. The default value is 16311 for the HTTPS protocol and 16310 for the HTTP protocol. The valid port values are from 1 to 65535 inclusive. This parameter is not required.
Polling Interval	How often the local data store is updated from the Authorization Policy Server by the policy client running on the portal server. The default is 30 minutes. Valid values are from 5 to 1440 minutes inclusive. This parameter is not required.

Table 41. Configuration information for the Authorization Policy Server (continued)

Field	Description
Policy Store Expiration Interval	<p>If the policy store cannot be updated from the Authorization Policy Server, this interval is the amount of time the local policy store will continue to be utilized from the last update. If the Authorization Policy Server cannot be accessed for the time interval specified by this parameter, all subsequent requests for dashboard data will fail with an authorization error until the Authorization Policy Server is available again. The default is 7 days and 0 hours. The value specified for hours must be in the range of 0-23 hours. If the expiration interval is set to 0 days and 0 hours, the policy store will never expire.</p> <p>This parameter is not required.</p>
User ID	<p>The name of the user that the portal server will use to access the IBM Dashboard Application Services Hub with Authorization Policy Server. This user must be added to the PolicyDistributor authorization policy core role or to a custom role that has been granted permission to perform the distribute operation for the role object type. See the <i>Using role-based authorization policies</i> chapter in the <i>IBM Tivoli Monitoring Administrator's Guide</i> and the <i>IBM Tivoli Monitoring Command Reference</i> for more details on creating authorization policy roles and assigning users to them.</p> <p>This parameter is required.</p>
Password	<p>The password for the user.</p> <p>This parameter is required.</p>
Confirm Password	<p>Confirm the password by entering it again.</p> <p>This parameter is required.</p>

Enter the required information for the Authorization Policy Server in the fields provided and click **OK**. For more information about the Tivoli Authorization Policy Server and Tivoli Authorization Policy Command-Line Interface, see “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318.

20. A message is displayed asking if you want to reconfigure the warehouse connection for the portal server.




Figure 41. Reconfigure warehouse connection information for the portal server

Do one of the following:

- Click **No** if you have not set up a Tivoli Data Warehouse.

Follow the instructions later in this book for implementing a Tivoli Data Warehouse solution, beginning with Chapter 19, “Tivoli Data Warehouse solutions,” on page 509. These instructions will direct you to reconfigure the connection between the portal server and the warehouse database after you have completed all preliminary setup tasks.

If you select **No**, go to Step 27 on page 240.

- Click **Yes** if you have completed the tasks for setting up a Tivoli Data Warehouse and want to configure the connection between the portal server and the Tivoli Data Warehouse database at this time. (You can choose to configure the connection later.) The prerequisite tasks and the information you need to configure the connection for each database type are described in the following steps.

For additional information about warehouse configuration, press the **Help** button.

21. Specify the database type to be used for the Warehouse Proxy data source.

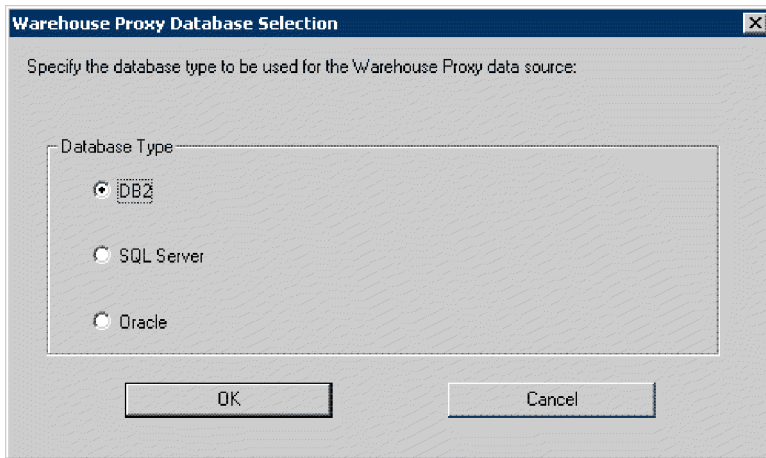


Figure 42. Warehouse Proxy Database Selection

If you select Oracle without having installed an Oracle ODBC driver an error message is displayed saying *Oracle ODBC driver not installed on this machine* and you cannot continue with the configuration. You can obtain Oracle ODBC drivers from the following Web site:
<http://www.oracle.com/technology/software/tech/windows/odbc/htdocs/utlsoft.html>

22. A window is displayed for you to configure the connection between the portal server and the *portal server database* (TEPS database). There is check box on the configuration window: **The database and the Warehouse Data Source do not exist and should be created by the Installer**. If you select this check box the installation program uses the information on this window to automatically perform the following tasks:
- Create the portal server database.
 - Create a database user for the portal server to use to access the database.
 - Configure the ODBC connection between the portal server and the database.

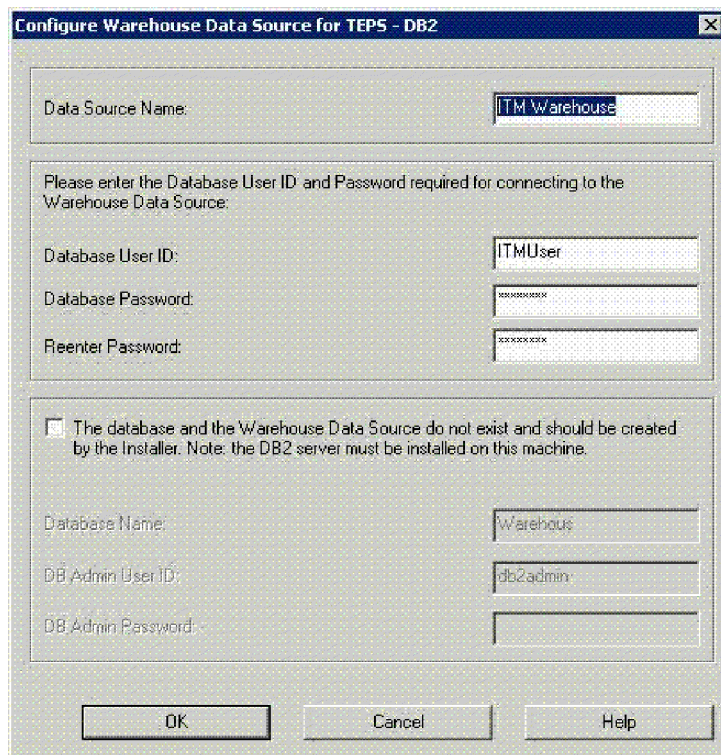


Figure 43. Configuration window for the portal server database using DB2 for Linux, UNIX, and Windows

Figure 43 shows the configuration window for a portal server database using DB2 for Linux, UNIX, and Windows. The configuration window for a Microsoft SQL Server database is similar. The fields on the configuration window are described in the following table:

Table 42. Configuration information for the portal server database

Field	DB2 for Linux, UNIX, and Windows default	MS SQL default	Description
Data Source Name	ITM Warehouse	ITM Warehouse	The name of the data source.
Database User ID	ITMUser	ITMUser	The name of the Windows OS user that the portal server will use to access the Tivoli Data Warehouse database.
Database Password	itmpswd1	itmpswd1	The password for the Windows user. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Reenter Password	itmpswd1	itmpswd1	Confirm the password by entering it again.

23. To configure a connection between the portal server and a DB2 for Linux, UNIX, and Windows data warehouse, complete the following steps:

a. Verify that you have completed the following tasks:

- Created a warehouse database using DB2 for Linux, UNIX, and Windows
- Created a warehouse user on the computer where you created the warehouse database.

The *warehouse user* is the user account (user ID and password) used by the portal server and other warehousing components to access the warehouse database.

- Activated the UNIX listeners on the computer where the warehouse database is located if the warehouse database is installed on UNIX systems.
- Installed a DB2 for Linux, UNIX, and Windows client on the portal server if the data warehouse is remote and the portal server database does not use DB2 for Linux, UNIX, and Windows.
- Cataloged the warehouse database on the computer where you are installing the portal server if the warehouse database is remote from the portal server.

These tasks are described in Chapter 23, “Tivoli Data Warehouse solution using DB2 for Linux, UNIX, and Windows,” on page 539.

- Gather the following information: data source name, database name, database administrator ID and password, warehouse user ID and password. The warehouse user ID is the one declared in the configuration panel of the Warehouse Proxy Agent. This user ID serves as the first part of the name of all the tables created in the Warehouse database. If you do not declare the same user ID when configuring the Tivoli Enterprise Portal server you will not be able to see the Warehouse tables with the portal client even if they exist in the database.
- Complete the “Configuring a Windows portal server (ODBC connection)” on page 559 procedure, starting from Step 4 on page 560.

For additional information about these parameters, press the **Help** button.

- To configure a connection between the portal server and a Microsoft SQL Server data warehouse, complete the following steps:

- Verify that you have completed the following tasks:
 - Created a warehouse database using Microsoft SQL Server.
 - Created a warehouse user on the computer where you created the warehouse database.
The *warehouse user* is the user account (user ID and password) used by the portal server and other warehousing components to access the warehouse database.
 - Installed a Microsoft SQL Server client on the portal server if the data warehouse is remote and the portal server database does not use Microsoft SQL Server.
 - Configured a remote client connection on the computer where you are installing the portal server if the warehouse database is remote from the portal server.

These tasks are described in Chapter 25, “Tivoli Data Warehouse solution using Microsoft SQL Server,” on page 595.

- Gather the following information: data source name, database name, database administrator ID and password, warehouse user ID and password. The warehouse user ID is the one declared in the configuration panel of the Warehouse Proxy Agent. This user ID serves as the first part of the name of all the tables created in the Warehouse database. If you do not declare the same user ID when configuring the Tivoli Enterprise Portal server you will not be able to see the Warehouse tables with the portal client even if they exist in the database.
- Complete the procedure “Configuring the portal server (ODBC connection)” on page 610, starting from Step 4 on page 610.

For additional information about these parameters, press the **Help** button.

- To configure a connection between the portal server and an Oracle data warehouse, complete the following steps:

- Verify that you have completed the following tasks:
 - Created a warehouse database using Oracle
 - Created a warehouse user on the computer where you created the warehouse database.
The *warehouse user* is the user account (user ID and password) used by the portal server and other warehousing components to access the warehouse database.
 - Activated the Oracle listener on the computer where the warehouse database is located
 - Installed an Oracle client on the portal server

- Created a TNS Service Name on the computer where you are installing the portal server if the warehouse database is remote from the portal server.

These tasks are described in Chapter 26, “Tivoli Data Warehouse solution using Oracle,” on page 619.

- Gather the following information: the data source name, database name, database administrator ID and password, and the warehouse user ID and password. The warehouse user ID is the one declared in the configuration panel of the Warehouse Proxy Agent. This user ID serves as the first part of the name of all the tables created in the Warehouse database. If you do not declare the same user ID when configuring the Tivoli Enterprise Portal server you will not be able to see the Warehouse tables with the portal client even if they exist in the database.
- Complete the procedure “Configuring a Windows portal server (ODBC connection)” on page 633, starting from Step 4 on page 633.

For additional information about these parameters, press the **Help** button.

26. Configure the default connector for the Common Event Console.




Figure 44. Common Event Console Configuration window

The default connector retrieves situation events reported to Tivoli Enterprise Monitoring Servers for display in the Common Event Console. You can configure connectors for other event management systems after you have completed the product installation. For configuration instructions, see the *IBM Tivoli Monitoring Administrator's Guide*.

Click **OK** to accept the default values or specify values for the following fields, then click **OK**:

Enable this connector

Select Yes to enable the connector to collect and display situation events in the Common Event Console, or No to configure the connector without enabling it. The connector is enabled by default.

Connector name

The name that is to be displayed in the Common Event Console for this connector. The default name is ITM1.

Maximum number of events for this connector

The maximum number of events that are to be available in the Common Event Console for this connector. The default value is 100 events.

View closed events

Select **No** to display only active events in the Common Event Console for this connector. Select **Yes** to view both active and closed events. By default, only active events are displayed.

27. Configure the default communication between any monitoring agents installed on this computer and the hub Tivoli Enterprise Monitoring Server:
 - a. Click **OK** to accept the default communications protocol.
 - b. Ensure that the host name and port number of the Tivoli Enterprise Monitoring Server are correct. Click **OK**.

For additional information about these parameters, press the **Help** button.

28. Click **Finish** to complete the installation.

Upgrading a 32 bit Windows portal server to a new system running 64 bit Windows

The installation process does not convert an existing 32 bit Tivoli Enterprise Portal Server running on Windows to 64 bit. However, you can migrate the portal server database onto a new (separate) 64 bit Windows system.

About this task

The following procedure allows you to migrate the portal server database onto a new 64 bit computer.

Procedure

1. Login to the system where the 32 bit Windows portal server is installed.
 - a. Open a command prompt and change to the following directory:
C:\IBM\ITM\CNPS
 - b. Run the following script to export the portal server database:
migrate-export.bat
 - c. Verify that the following file was generated, and copy it to a safe location:
C:\IBM\ITM\CNPS\SQLLIB\saveexport.sql
2. Login to the new 64 bit Windows system.
 - a. Install the portal server to C:\IBM\ITM\CNPS.
For more information, see “Windows: Installing the portal server” on page 227.
 - b. Install application support for base agents.
For more information, see “Installing and enabling application support” on page 275.
 - c. Install application support for any additional agents that are not part of the base installation.
For more information, see “Configuring application support for nonbase monitoring agents” on page 279.
 - d. Copy the saveexport.sql file from step 1 c into C:\IBM\ITM\CNPS on the new 64 bit Windows system.
 - e. Open a command prompt and change to the following directory:
C:\IBM\ITM\CNPS
 - f. Run the following script to import the portal server database:
migrate-import.bat

- g. Run the following script to reapply the latest application support:
`buildpresentation.bat`
- h. Restart the Tivoli Enterprise Portal Server.

Linux or AIX: Installing the portal server

Complete the procedures in this section to install and configure the Tivoli Enterprise Portal Server and portal client on a Linux or AIX computer.

Important: Run these installation and configuration procedures as either the root user or as the DB2 for Linux, UNIX, and Windows administrator. If you are configuring the Tivoli Enterprise Portal Server using the DB2 for Linux, UNIX, and Windows administrator ID, then:

- The configuration ID must be the same as the ID used to install the Tivoli Enterprise Portal Server.
- The configuration ID must be the same as the ID specified during the configuration dialog for the DB2 Admin ID.
- The configuration ID must have the proper DB2 for Linux, UNIX, and Windows authority/rights to attach to the DB2 Instance Name specified during the configuration dialog.
- The ID specified during the configuration dialog for the DB2 User ID must be an already existing ID. You may not allow the configuration dialog to attempt to create the user. For the GUI dialog, ensure that the check box for **Create the user** is cleared. For the CLI dialog, ensure that the response for Create New User is NO. After you have installed and configured the portal server, you can use a different user to run the portal server, as long as that user has access to the binaries used by the portal server.

Note: Security policies for newly created users usually auto-expire the password after the first use and require the user to set a new password as part of the initial logon process. Until this is done, the installer always fails because the user password is expired and must be reset. Before running the installer, the user must invoke a process like ssh or telnet to log onto the target user ID and set the password appropriately.

Table 43. Steps for installing a portal server on a Linux or AIX computer

Steps	Where to find information
Install the portal server.	"Installing the portal server on Linux or AIX"
Configure the portal server.	"Configuring the portal server on Linux or AIX: command-line procedure" on page 243 –OR– "Configuring the portal server on Linux or AIX: GUI procedure" on page 250
Start the portal server.	"Starting the portal server" on page 260

Prerequisites for users installing on Linux on zSeries

If you plan to install the Tivoli Enterprise Portal Server on Linux for zSeries, make sure you complete these steps before beginning your installation.

- Ensure your Linux environment provides at least 10 gigabytes of free disk space.
- Turn off the Just-In-Time Java compiler.
- Users of RHEL version 5: set security-enhanced Linux (SELinux) to **permissive**.

Installing the portal server on Linux or AIX

Complete the following steps to install the portal server.

Before you begin

References to “browser client” in this section means both the browser and Java webstart clients, as both client types are installed as a single unit. These components are not optional when you install the portal server. The desktop client is a separately installed component that is optional.

Procedure

1. In the directory where you extracted the installation files, run the following command:
`./install.sh`
2. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory (/opt/IBM/ITM) or type the full path to a different directory.

Note:

- a. If you specify an incorrect directory name, you will receive the following error:

The IBM Tivoli Monitoring installation directory cannot exceed 80 characters or contain non-ASCII, special or double-byte characters.

The directory name can contain only these characters:

"abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ _\.:0123456789()~-./".

- b. You must not specify the path of the directory containing `./install.sh` as your IBM Tivoli Monitoring home directory. On certain platforms, this can cause the plug-in JAR files to overwrite themselves and become zero length files. The installation will fail as a result.
3. If the installation directory does not already exist, you are asked if you want to create it. Type y to create this directory and press Enter.
 4. The following prompt is displayed:
Select one of the following:
1) Install products to the local host.
2) Install products to depot for remote deployment (requires TEMS).
3) Install TEMS support for remote seeding
4) Exit install.

Please enter a valid number:

Type 1 to start the installation and display the software license agreement.

5. Press Enter to read through the agreement.
6. Type 1 to accept the agreement and press Enter.
7. Enter a 32-character encryption key or press Enter to accept the default key. This key should be the one used during the installation of the monitoring server to which this portal server will connect.
A numbered list of available operating systems is displayed.
8. Type 4 to install the portal server for your current operating system. Press Enter.
A message is displayed indicating that the Tivoli Enterprise Portal Server is about to be installed.

Note: The Eclipse Help Server is automatically installed when you install the Tivoli Enterprise Portal Server.

9. Type 1 to confirm the installation.
The installation begins.
10. After the Tivoli Enterprise Portal Server is installed, you are asked whether you want to install additional products or product support packages. Type 1 and press Enter.
The installer presents a numbered list of products and application support packages.
11. Install the required application support packages.

All monitoring agents require that application support files be installed on the monitoring servers (hub and remote), portal server, and portal desktop clients in your environment. Application support files contain the information required for agent-specific workspaces, helps, predefined situations, and other data.

This step installs the application support files for base monitoring agents. The *base* monitoring agents are included with the base IBM Tivoli Monitoring installation package. For detailed information about application support, see “Installing and enabling application support” on page 275.

When you entered 1 in the preceding step, the installer presents a numbered list of items, including the following application support packages:

- Tivoli Enterprise Portal Browser Client support
- Tivoli Enterprise Portal Server support

Note: The Tivoli Enterprise Portal Browser Client support package is portal server code that supports the browser clients. You must install the browser client support package on the computer where you install the portal server if you want to connect to it using a browser client.

Complete the following steps to install the portal server and browser client support packages for the base monitoring agents:

- a. Type the number that corresponds to Tivoli Enterprise Portal Browser Client support and press Enter.
A numbered list of base monitoring agents is displayed.
- b. Type the numbers that correspond to the base monitoring agents for which you want to install the application support package, or type the number that corresponds to All of the above. Type the numbers on the same line separated by spaces or commas (.). Press Enter.
It is best to select all of the base monitoring agents (All of the above) so you do not need to reconfigure application support as new agent types are added to your environment.
- c. Type 1 to confirm the installation and press Enter.
The installation begins.
- d. After the support package is installed, you are asked whether you want to install additional products or product support packages. Enter 1 and repeat the preceding steps for the Tivoli Enterprise Portal Server support package.

Note: This step installs the application support files. However, you must *enable* the application support by configuring the portal server. The next two sections show you how to configure the portal server.

12. After you are finished installing the portal server and browser client packages, you are asked whether you want to install additional products or product support packages. Type 2 and press Enter.
13. If your IBM Tivoli Monitoring environment is not already secured you will be asked at this point if you want to secure it. If your IBM Tivoli Monitoring environment is already secured this question is skipped. The product installation process creates the majority of directories and files with world write permissions. IBM Tivoli Monitoring provides the secureMain utility to help you keep the monitoring environment secure. You can secure your installation now, or manually execute the secureMain utility later. For more information, see Appendix G, “Securing your IBM Tivoli Monitoring installation on Linux or UNIX,” on page 925.

Configuring the portal server on Linux or AIX: command-line procedure

Complete the following steps to configure the Tivoli Enterprise Portal Server from the command-line on Linux or AIX.

About this task

Configure the portal server using either the command-line procedure in this section or the GUI procedure in “Configuring the portal server on Linux or AIX: GUI procedure” on page 250.

Either of these configuration procedures accomplishes the following tasks:

- Automatically enables application support on the portal server for the base monitoring agents. (See Step 11 on page 242 of the installation procedure.)
- Includes steps for configuring the connection between the portal server and the following components:
 - The hub monitoring server
 - The portal server database
 - The Tivoli Data Warehouse database
 - Tivoli Authorization Policy Server

Note: If you have not set up the Tivoli Data Warehouse, complete this procedure but accept the default values at the prompts for configuring the connection to the data warehouse. You can reconfigure the connection after you set up the warehouse. See Step 9 on page 246 for more information.

Procedure

1. Log on to the computer where the Tivoli Enterprise Portal Server is installed.
2. At the command-line change to the *ITMinstall_dir/bin* directory, where *ITMinstall_dir* is the directory where you installed the product.
3. Run the following command to start configuring the Tivoli Enterprise Portal Server:

```
./itmcmd config -A cq
```

where *cq* is the product code for the portal server.

4. Edit the ITM Connector settings.

The ITM Connector retrieves situation events reported to Tivoli Enterprise Monitoring Servers for display in the Common Event Console. If you do not enable the connector, you will not see any events in the Common Event Console. You can configure connectors for other event management systems after you have completed the product installation. For configuration instructions, see the *IBM Tivoli Monitoring Administrator's Guide*.

To enable the connector:

- a. Press Enter to accept the default value for the following prompt:
Edit 'ITM Connector' settings? [1=Yes, 2=No] (default is: 1):
 - b. Press Enter to enable the connector.
 - c. Press Enter to accept the default name of ITM1 or type your preferred name and press Enter. This is the name that is to be displayed in the Common Event Console for this connector.
 - d. Press Enter to accept the default number of events (100) that are to be available in the Common Event Console for this connector, or type the number of events you would like to see displayed and press Enter.
 - e. Type 2 and press Enter to display only active events in the Common Event Console for this connector. Type 1 and press Enter to view both active and closed events. By default, only active events are displayed.
 - f. Type 2 and press Enter to skip defining data for extra columns in the Common Event Console.
When you define a Tivoli Enterprise Console or Tivoli Netcool/OMNIbus connector, you can define the information that is to be mapped to each of these customizable columns. See the *IBM Tivoli Monitoring Administrator's Guide* for information on configuring these connectors.
5. Press Enter when you are asked if the agent connects to a monitoring server. (Although the prompt refers to an *agent*, this command is used to configure the portal server.)
 6. Configure the connection between the portal server and the hub monitoring server:
 - a. Type the host name for the hub monitoring server and press Enter.
 - b. Type the protocol that the hub monitoring server uses to communicate with the portal server. You have the following choices: IP.PIPE, IP.SPIPE, IP.UDP, or SNA. The IP.UDP and SNA

protocols have been deprecated in favour of IP.PIPE and IP.SPIPE. You should avoid using the IP.UDP and SNA protocols even though they are available configuration options.

- c. If you want to set up a backup protocol, enter that protocol and press Enter. If you do not want to use a backup protocol, press Enter without specifying a protocol.
- d. Depending on the type of protocol you specified, provide the following information as described in Table 44 when prompted:

Table 44. Hub monitoring server protocols and values

Protocol	Value	Definition
IP.UDP (deprecated, use IP.PIPE or IP.SPIPE instead)	IP Port Number	The port number for the monitoring server. The default is 1918.
SNA (deprecated, use IP.PIPE or IP.SPIPE instead)	Net Name	The SNA network identifier for your location.
	LU Name	The LU name for the monitoring server. This LU name corresponds to the Local LU Alias in your SNA communications software.
	Log Mode	The name of the LU6.2 LOGMODE. The default value is CANCTDCS.
IP.PIPE	IP.PIPE Port Number	The port number for the monitoring server. The default is 1918.
IP.SPIPE	IP.SPIPE Port Number	The port number for the monitoring server. The default is 3660.

- e. Press Enter when you are asked if you want to configure the connection to a secondary monitoring server. The default value is none.
 - f. Press Enter to accept the default value for the Optional Primary Network Name (none).
 - g. Press Enter to accept the default setting for SSL between the portal server and clients (N). By default, SSL is disabled. To enable SSL, type 1 and press Enter.
7. Configure the connection between the portal server and the portal server database:
- a. Type 1 if your site is using the embedded Derby database, 2 if you're using DB2 for Linux, UNIX, and Windows.
 - b. Type the DB2 for Linux, UNIX, and Windows instance name. The default value is db2inst1. Press Enter.
 - c. Type the DB2 for Linux, UNIX, and Windows administrator ID. The default value is db2inst1. Press Enter.
- Note:** The DB2 for Linux, UNIX, and Windows administrator account was created during DB2 for Linux, UNIX, and Windows installation.
- d. Type the password for the DB2 for Linux, UNIX, and Windows administrator ID, and press Enter.
 - e. Confirm the password for the DB2 for Linux, UNIX, and Windows administrator ID by typing it again. Press Enter.
8. If you are configuring DB2 for Linux, UNIX, and Windows for the portal server (instead of the embedded Derby database), complete the following parameters as well:
- a. Type the name of the portal server database. The default value is TEPS. Press Enter.
 - b. Type the login name of the database user that the portal server will use to access the database. The default value is itmuser. Press Enter.
 - c. Type the password for the database user and press Enter.
 - d. Confirm the password for the database user by typing it again. Press Enter.

- e. You are asked if you want to create the DB2 for Linux, UNIX, and Windows login user if it does not exist. Type 1 and press Enter.
9. You are asked for the database parameters for either DB2 for Linux, UNIX, and Windows or Oracle for the Tivoli Data Warehouse. Enter D for DB2 for Linux, UNIX, and Windows, J for Oracle (JDBC). DB2 for Linux, UNIX, and Windows is the default.

Note: This prompt and all remaining prompts ask for information to configure the connection between the portal server and the Tivoli Data Warehouse database. If you have not set up a Tivoli Data Warehouse, accept the default values at these prompts. Follow the instructions later in this book for implementing a Tivoli Data Warehouse solution, beginning with Chapter 19, “Tivoli Data Warehouse solutions,” on page 509. These instructions will direct you to reconfigure the connection between the portal server and the warehouse database after you have completed all preliminary setup tasks.

10. Do one of the following:
 - If you have not set up a Tivoli Data Warehouse:
 - a. Press Enter to accept the DB2 for Linux, UNIX, and Windows default (even if you are going to create the Tivoli Data Warehouse using Oracle).
 - b. Press Enter at all remaining prompts to accept the default values.

or

- Configure a connection between the portal server and a DB2 for Linux, UNIX, and Windows warehouse database.

Before you configure the connection, verify that you have already completed the following tasks:

- Created a warehouse database using DB2 for Linux, UNIX, and Windows
- Created a warehouse user on the computer where you created the warehouse database.
The *warehouse user* is the user account (user ID and password) used by the portal server and other warehousing components to access the warehouse database.
- Cataloged the warehouse database on the computer where you are installing the portal server if the warehouse database is remote from the portal server
- Activated the UNIX listeners on the computer where the warehouse database is located if the warehouse database is installed on UNIX.

These tasks are described in Chapter 23, “Tivoli Data Warehouse solution using DB2 for Linux, UNIX, and Windows,” on page 539.

To perform the steps for configuring the connection between the portal server and the DB2 for Linux, UNIX, and Windows data warehouse, you will need the name of the warehouse database and the user ID and password of the warehouse user. Press Enter after answering each prompt:

- a. Press Enter to accept the DB2 for Linux, UNIX, and Windows default.
- b. Enter the name of the Tivoli Data Warehouse database. The default value is WAREHOUS.
- c. Enter the user ID of the warehouse user. The default value is itmuser.

Note: The warehouse user ID must be the one declared in the configuration panel of the Warehouse Proxy Agent. This user ID serves as the first part of the name of all the tables created in the Warehouse database. If you do not declare the same user ID when configuring the Tivoli Enterprise Portal server you will not be able to see the Warehouse tables with the portal client even if they exist in the database.

- d. Enter the password of the warehouse user.
- e. Confirm the password by entering it again.

–OR–

- Configure a connection between the portal server and an Oracle warehouse database.

Before you configure the connection, verify that you have already completed the following tasks:

- Created a warehouse database using Oracle

- Created a warehouse user on the computer where you created the warehouse database.
The *warehouse user* is the user account (user ID and password) used by the portal server and other warehousing components to access the warehouse database.
- Activated the Oracle listener on the computer where the warehouse database is located.
- Installed an Oracle JDBC Type 4 driver on the portal server.

These tasks are described in Chapter 26, “Tivoli Data Warehouse solution using Oracle,” on page 619.

To perform the steps for configuring the connection between the portal server and the Oracle data warehouse, you will need:

- The name of the warehouse database and the user ID and password of the warehouse user
- The location, name, and URL of the JDBC driver

Press Enter after answering each prompt:

- a. Enter `oracle` at the prompt asking if you are using DB2 for Linux, UNIX, and Windows or Oracle for the Warehouse.
- b. Enter the name of the Tivoli Data Warehouse database. The default value is `WAREHOU`.
- c. Enter the user ID of the warehouse user. The default value is `itmuser`.

Note: The warehouse user ID must be the one declared in the configuration panel of the Warehouse Proxy Agent. This user ID serves as the first part of the name of all the tables created in the Warehouse database. If you do not declare the same user ID when configuring the Tivoli Enterprise Portal server you will not be able to see the Warehouse tables with the portal client even if they exist in the database.

- d. Enter the password of the warehouse user.
- e. Confirm the password by entering it again.
- f. Enter the full path name of the Oracle Type 4 JDBC driver JAR file as follows:
`oracleinstalldir/jdbc/lib/ojdbc14.jar`

where `oracleinstalldir` is the directory location of the JDBC driver JAR file on this computer.

- g. Enter the following JDBC driver name:
`oracle.jdbc.driver.OracleDriver`
 - h. Enter the JDBC driver URL. This is the Oracle-defined URL that identifies the locally or remotely installed Oracle instance used for the Tivoli Data Warehouse. The following entry is an example:
`jdbc:oracle:thin:@localhost:1521:WAREHOU`
 - If the warehouse database is located on a remote computer, replace `localhost` with the host name of the remote computer.
 - Change the default port number (1521) and Tivoli Data Warehouse name (WAREHOU) if they are different.
 - i. Enter any user-defined attributes that are used to customize the behavior of the driver connection. Use semi-colons (;) to delimit the attributes. Press Enter to finish the configuration.
11. Configure LDAP Security. By default, **LDAP Security: Validate User with LDAP** is enabled (1=Yes, 2=No). Press Enter if you want the portal server to use an LDAP registry to authenticate users and, optionally, to enable single sign on.

A best practice is to select 2 to install and configure the portal server with LDAP validation disabled. After a basic installation of IBM Tivoli Monitoring has been completed and tested then you can reconfigure the portal server and enable LDAP user validation using the instructions in the *Enabling user authentication* chapter of the *IBM Tivoli Monitoring Administrator's Guide*.

Note: If you enable LDAP user validation, the LDAP server must be started so that Tivoli Monitoring can verify a connection to the LDAP server. If a connection cannot be established, LDAP user validation will be disabled when the portal server installation completes.

12. If you select 1 for enable LDAP user validation, you are prompted for the following LDAP configuration parameters:

- LDAP base:
- LDAP DN Base Entry: (Default is: o=ITMSSOEntry)
- LDAP bind ID:
- LDAP bind password:
- Re-type: LDAP bind password:
- LDAP Port number: (Default is: 389)
- LDAP host name: (Default is: localhost)

For more information about configuring these fields, see *Prerequisites for configuring LDAP authentication on the portal server* in the *IBM Tivoli Monitoring Administrator's Guide*.

13. Configure Single Sign On if you enabled LDAP user validation. By default, Enable Single Sign On is disabled (1=Yes, 2=No). Type 1 if you want the portal server to use single signon with other applications such as IBM Dashboard Application Services Hub or with other applications that launch the portal client.

A best practice is to install and configure the portal server with single signon disabled. After a basic installation of IBM Tivoli Monitoring has been completed and tested, you can reconfigure the portal server and enable single signon using the instructions in the *Enabling user authentication* chapter of the *IBM Tivoli Monitoring Administrator's Guide*.

- a. If you selected 1 to enable Single Sign On, you are prompted to enter values for the following parameters:

- Domain name
- Realm name

For more information about configuring these fields and setting up SSO, see *Prerequisites for configuring LDAP authentication on the portal server* and *About single sign-on* in the *IBM Tivoli Monitoring Administrator's Guide*.

14. Configure dashboard data provider. By default, the **dashboard data provider** is disabled (1=Yes, 2=No). Type 1 if Dashboard Application Services Hub will be used with monitoring dashboard applications such as IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, or IBM Infrastructure Management Capacity Planner for PowerVM, or will be used to create or enable custom dashboard pages for monitoring agents. (See the *Preparing your dashboard environment* topic in the *IBM Tivoli Monitoring Administrator's Guide* for details on additional configuration tasks to perform after installation.)

- The domain override value is optional and is only valid when the dashboard data provider is enabled. It changes the default dashboard data provider ID and domain name for authorization policies from itm.<Hub monitoring server name> to itm.<domain override value>. The value may not exceed 124 characters. You should configure a domain override value for these scenarios:
 - The Hot Standby high availability feature is being used for the Hub monitoring server. By configuring a domain override value, the dashboard data provider ID and domain name will not change when the portal server is configured to connect to the new acting Hub monitoring server. If you do not configure a domain override value in this scenario, you must reconfigure the connection between the IBM Dashboard Application Services Hub and the dashboard data provider and update any domain specific authorization policies when the portal server is configured to connect to the new acting Hub monitoring server.
 - You have multiple Hub monitoring servers that are using a common set of authorization policies for controlling dashboard access and you want to create some domain specific

authorization policies. You should specify a domain override value for this scenario if you want to use a more user friendly domain name in your authorization policies than the default value of itm.<Hub monitoring server name>.

- If you modify the domain override after you have configured a connection in your Dashboard Application Services Hub Server to the dashboard data provider then you must delete the connection and re-add it. See the *IBM Tivoli Monitoring Administrator's Guide* for details on how to configure a dashboard data provider connection. Also if you have created any domain specific authorization policies using the default domain name then you must delete the permissions that use the previous domain name and create new permissions that use the new domain name when you change the domain override value. See the *IBM Tivoli Monitoring Command Reference* for details on the tivcmd CLI commands used to create and work with authorization policies.
15. If you selected Yes in step 14 on page 248 you are asked if you want to enable authorization policies (1=Yes, 2=No. Default is 1). Enable authorization policies if you want to use authorization policies to control which managed systems and managed system groups a user can access in monitoring dashboards of the IBM Dashboard Application Services Hub. Once authorization policies are enabled, only authorized users will see data in monitoring dashboards. You should only enable authorization policies if your administrators have already created the initial set of policies for dashboard user access using the Tivoli Authorization Policy Server and tivmd Command Line Interface for Authorization Policy. See the *Using role-based authorization policies* topic in the *IBM Tivoli Monitoring Administrator's Guide* for details on how to create and work with authorization policies.

Note: If authorization policies have not been created yet, do not enable authorization policies. After the authorization policies are ready reconfigure the portal server and enable authorization policies. The configuration parameters are described in the following table:

Table 45. Configuration information for the Authorization Policy Server

Parameter	Description
Hostname or IP Address	IP Address or fully qualified hostname of the IBM Dashboard Application Services Hub with the Authorization Policy Server. This parameter is required.
Protocol	Choose the protocol used to connect to the IBM Dashboard Application Services Hub with the Authorization Policy Server. The default value is HTTPS. This parameter is not required.
Port	Choose the port used to connect to the IBM Dashboard Application Services Hub with the Authorization Policy Server. The default value is 16311 for the HTTPS protocol and 16310 for the HTTP protocol. The valid port values are from 1 to 65535 inclusive. This parameter is not required.
Polling Interval	How often the local data store is updated from the Authorization Policy Server by the policy client running on the portal server. The default is 30 minutes. Valid values are from 5 to 1440 minutes inclusive. This parameter is not required.

Table 45. Configuration information for the Authorization Policy Server (continued)

Parameter	Description
Policy Store Expiration Interval	<p>If the policy store cannot be updated from the Authorization Policy Server, this interval is the amount of time the local policy store will continue to be utilized from the last update. If the Authorization Policy Server cannot be accessed for the time interval specified by this parameter, all subsequent requests for dashboard data will fail with an authorization error until the Authorization Policy Server is available again. The default is 7 days and 0 hours. The value specified for hours must be in the range of 0-23 hours. If the expiration interval is set to 0 days and 0 hours, the policy store will never expire.</p> <p>This parameter is not required.</p>
User ID	<p>The name of the user that the portal server will use to access the IBM Dashboard Application Services Hub with Authorization Policy Server. This user must be added to the PolicyDistributor authorization policy core role or to a custom role that has been granted permission to perform the distribute operation for the role object type. See the <i>Using role-based authorization policies</i> chapter in the <i>IBM Tivoli Monitoring Administrator's Guide</i> and the <i>IBM Tivoli Monitoring Command Reference</i> for more details on creating authorization policy roles and assigning users to them.</p> <p>This parameter is required.</p>
Password	<p>The password for the user.</p> <p>This parameter is required.</p>
Confirm Password	<p>Confirm the password by entering it again.</p> <p>This parameter is required.</p>

Results

A message is displayed telling you that InstallPresentation is running, and then a message telling you that the installation has completed.

Configuring the portal server on Linux or AIX: GUI procedure

Complete the following steps to configure the Tivoli Enterprise Portal Server from the Tivoli Enterprise Monitoring Services window on Linux or AIX.

About this task

Configure the portal server using either the GUI procedure in this section or the command-line procedure in “Configuring the portal server on Linux or AIX: command-line procedure” on page 243.

Either of these configuration procedures accomplishes the following tasks:

- Automatically enables application support on the portal server for the base monitoring agents.
- Includes steps for configuring the connection between the portal server and the following components:
 - The hub monitoring server
 - The portal server database
 - The Tivoli Data Warehouse database
 - Tivoli Authorization Policy Server

Note: If you have not set up the Tivoli Data Warehouse, complete this procedure but accept the default values at the prompts for configuring the connection to the data warehouse. You can reconfigure the connection after you set up the warehouse. See Step 9 on page 246 for more information.

Procedure

- 1. Log on to the computer where the Tivoli Enterprise Portal Server is installed.
- 2. Start the Manage Tivoli Enterprise Monitoring Services utility:
 - a. Change to the bin directory:
`cd install_dir/bin`
 - b. Run the following command by using the parameters described in the next table:
`./itmcmd manage [-h ITMinstall_dir]`

where:

-h	(optional) An option used to specify the installation directory.
ITMinstall_dir	The directory where the portal server is installed. The default installation directory is /opt/IBM/ITM.

The Manage Tivoli Enterprise Monitoring Services window is displayed.

- 3. Right-click **Tivoli Enterprise Portal Server** and click **Configure**.
The Common Event Console Configuration window is displayed.

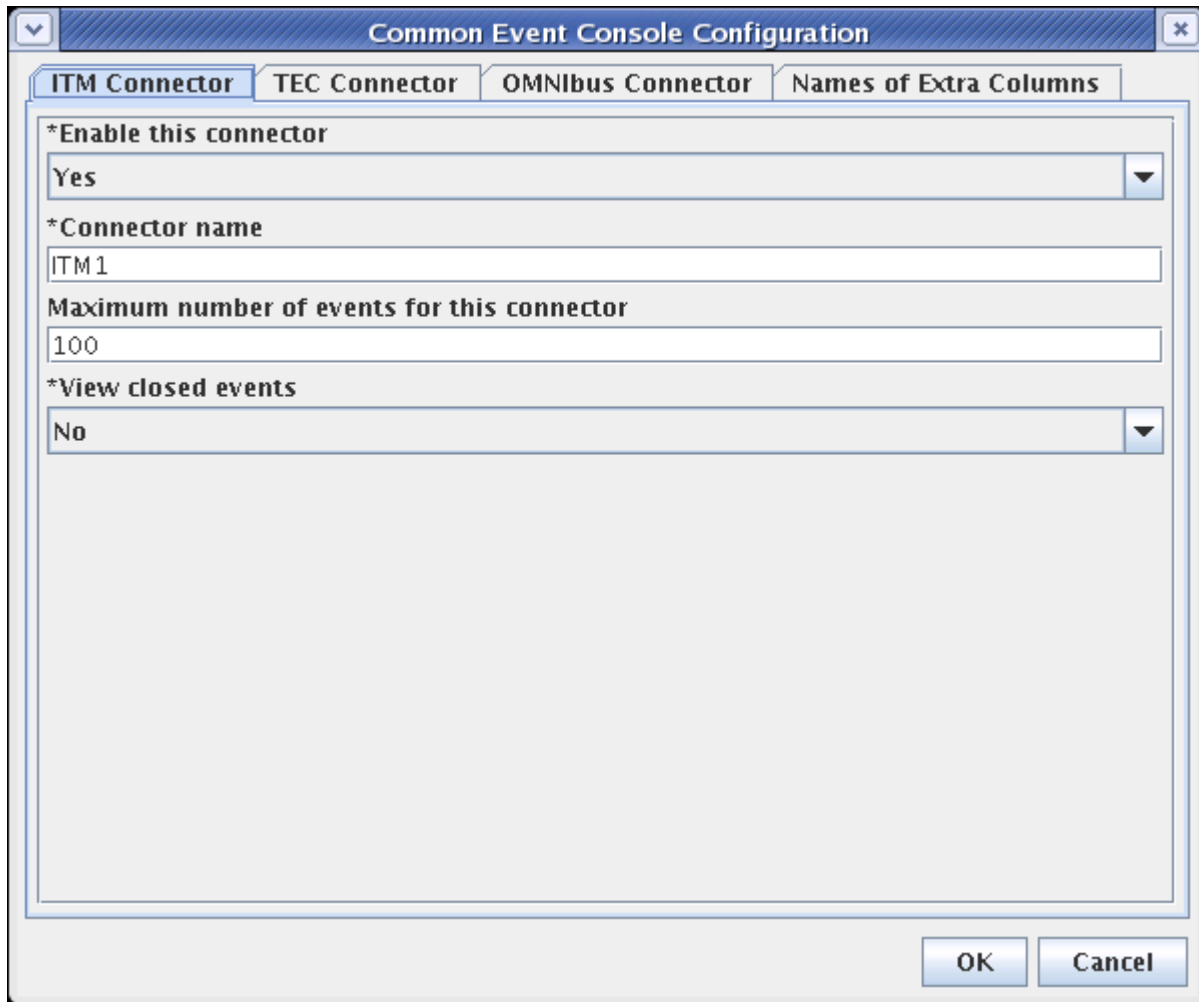


Figure 45. Common Event Console Configuration window

4. Click **OK** to accept the default values for the ITM Connector, or specify your preferred values and then click OK.

The ITM Connector retrieves situation events reported to Tivoli Enterprise Monitoring Servers for display in the Common Event Console. You can configure connectors for other event management systems after you have completed the product installation. For configuration instructions, see the *IBM Tivoli Monitoring Administrator's Guide*.

The Configure Tivoli Enterprise Portal Server window is displayed.

Figure 46. Registering the portal server with the Tivoli Enterprise Monitoring Server

5. On the TEMS Connection page, enter information about the Tivoli Enterprise Monitoring Server to which the Tivoli Enterprise Portal Server connects:
 - a. Enter the host name of the monitoring server in the **TEMS Hostname** field. (If the field is not active, clear the **No TEMS** check box.)
 - b. Select the communications protocol that the monitoring server uses from the **Protocol** drop-down list.
 - If you select SNA, enter information in the **Net Name**, **LU Name**, and **LOG Mode** fields.
 - If you select IP.UDP, IP.PIPE or IP.SPIPE, enter the port number of the monitoring server in the **Port Number** field.

For information about these fields, see Table 44 on page 245.

- c. Select **Validate user with LDAP** if you want the portal server to use an LDAP registry to authenticate users and, optionally, to enable single signon. A best practice is to install and configure the portal server with LDAP validation disabled. After a basic installation of IBM Tivoli Monitoring has been completed and tested, you can reconfigure the portal server and enable LDAP user validation using the instructions in the *Enabling user authentication* chapter of the *IBM Tivoli Monitoring Administrator's Guide*.

Note: If you enable LDAP user validation, the LDAP server must be started so that Tivoli Monitoring can verify a connection to the LDAP server. If a connection cannot be established, LDAP user validation will be disabled when the portal server installation completes.

- d. Select **Enable Single Sign On** if you enabled LDAP user validation and you want the portal server to use single signon with other applications such as IBM Dashboard Application Services Hub or with other applications that launch the portal client. A best practice is to install and configure the portal server with single signon disabled. After a basic installation of IBM Tivoli Monitoring has been completed and tested, you can reconfigure the portal server and enable single signon using the instructions in the *Enabling user authentication* chapter of the *IBM Tivoli Monitoring Administrator's Guide*.

6. Click the **Agent Parameters** tab.

Configure Tivoli Enterprise Portal Server

TEMS Connection | **Agent Parameters** | Dashboard data provider

TEPS DB | **Warehouse DB**

DB2 | DB2

TEPS database creation

DB2 instance name : db2inst1

DB2 admin ID : db2inst1

DB2 admin password :

Re-type DB2 admin password :

TEPS DB2 database name : TEPS

TEPS database connection parameters

TEPS DB user ID : itmuser

TEPS DB user password :

Re-type TEPS DB user password :

☒ Create TEPS DB user ID if not found?

Warehouse database connection parameters

Warehouse database name : WAREHOUS

Warehouse DB user ID : itmuser

Warehouse user password :

Re-type warehouse user password :

☐ Edit host specific configuration

Host :

Save | Reload | Help | Cancel

Figure 47. Configuring database connections for the portal server

7. Configure the connection between the portal server and the portal server database by entering information in the fields described in the following table:

Table 46. Configuration information for the Tivoli Enterprise Portal Server database

Field	Default value	Description
DB2 instance name	db2inst1	The DB2 for Linux, UNIX, and Windows instance name. Not required if the embedded Derby database is used for the portal server and Oracle is selected for the warehouse database.
DB2 admin ID	db2inst1	The DB2 for Linux, UNIX, and Windows administrator ID. The DB2 for Linux, UNIX, and Windows administrator account was created during DB2 for Linux, UNIX, and Windows installation. Not required if the embedded Derby database is used for the portal server and Oracle is selected for the warehouse database.
DB2 admin password	(no default)	The password for the DB2 for Linux, UNIX, and Windows administrator ID. Not required if the embedded Derby database is used for the portal server and Oracle is selected for the warehouse database.
Re-type DB2 admin password	(no default)	The password for the DB2 for Linux, UNIX, and Windows administrator ID. Not required if the embedded Derby database is used for the portal server and Oracle is selected for the warehouse database.
TEPS DB2 database name	TEPS	The Tivoli Enterprise Portal Server database name. Disabled if the embedded Derby database is chosen for the portal server database.
TEPS DB user ID	itmuser	The login name of the warehouse user that the portal server will use to access the database. Disabled if the embedded Derby database is chosen for the portal server database.
TEPS DB user password	(no default)	The password for the warehouse user ID. Disabled if the embedded Derby database is chosen for the portal server database.
Re-type TEPS DB user password	(no default)	The password for the warehouse user ID. Disabled if the embedded Derby database is chosen for the portal server database.
Create TEPS DB user ID if not found?	yes	This check box is selected by default. If the database login account (user ID and password) that you specified in the preceding fields does not exist, it is created. Disabled if the embedded Derby database is chosen for the portal server database.

- Optionally configure the connection between the portal server and the Tivoli Data Warehouse database.

Note: If you have not set up a Tivoli Data Warehouse, accept the default values for these fields. Follow the instructions later in this book for implementing a Tivoli Data Warehouse solution,

beginning with Chapter 19, “Tivoli Data Warehouse solutions,” on page 509. These instructions will direct you to reconfigure the connection between the portal server and the warehouse database after you have completed all preliminary setup tasks.

The bottom section of the **Agent Parameters** tab contains fields for configuring the connection between the portal server and a Tivoli Data Warehouse database using DB2 for Linux, UNIX, and Windows or Oracle. (See Figure 48.)

Figure 48. Configuration information for the Tivoli Data Warehouse using an Oracle database

Do one of the following:

- If you have not set up the Tivoli Data Warehouse and do not want to enable the dashboard data provider, click **Save** to save your settings and close the window.
–OR–
- Configure a connection between the portal server and a DB2 for Linux, UNIX, and Windows warehouse database.

Before you configure the connection, verify that you have already completed the following tasks:

- Created a warehouse database using DB2 for Linux, UNIX, and Windows

- Created a warehouse user on the computer where you created the warehouse database.
The *warehouse user* is the user account (user ID and password) used by the portal server and other warehousing components to access the warehouse database.
- Cataloged the warehouse database on the computer where you are installing the portal server if the warehouse database is remote from the portal server.
- Activated the UNIX listeners on the computer where the warehouse database is located if the warehouse database is installed on UNIX.

These tasks are described in Chapter 23, “Tivoli Data Warehouse solution using DB2 for Linux, UNIX, and Windows,” on page 539.

To configure the connection between the portal server and the DB2 for Linux, UNIX, and Windows data warehouse, you will need the name of the warehouse database and the user ID and password of the warehouse user.

To perform the configuration, complete the procedure “Configuring a Linux or AIX portal server (DB2 for Linux, UNIX and Windows CLI connection)” on page 561, starting from Step 4 on page 561.

–OR–

- Configure a connection between the portal server and an Oracle warehouse database.

Before you configure the connection, verify that you have already completed the following tasks:

- Created a warehouse database using Oracle
- Created a warehouse user on the computer where you created the warehouse database.
The *warehouse user* is the user account (user ID and password) used by the portal server and other warehousing components to access the warehouse database.
- Activated the Oracle listener on the computer where the warehouse database is located
- Installed an Oracle JDBC Type 4 driver on the portal server.

These tasks are described in Chapter 26, “Tivoli Data Warehouse solution using Oracle,” on page 619.

To configure the connection between the portal server and the Oracle data warehouse, you will need:

- The name of the warehouse database and the user ID and password of the warehouse user
- The location, name, and URL of the JDBC driver

To perform the configuration, complete the procedure “Configuring a Linux or AIX portal server (JDBC connection)” on page 634, starting from Step 4 on page 634.

9. Select the **Enable the dashboard data provider** check box if Dashboard Application Services Hub will be used with monitoring dashboard applications such as IBM Infrastructure Management Dashboards for Servers, IBM Infrastructure Management Dashboards for VMware, IBM Infrastructure Management Capacity Planner for VMware, IBM Infrastructure Management Capacity Planner for PowerVM, or will be used to create or enable custom dashboard pages for monitoring agents. (See the *Preparing your dashboard environment* topic in the *IBM Tivoli Monitoring Administrator's Guide* for details on additional configuration tasks to perform after installation.)
- The domain override value is optional and is only valid when the dashboard data provider is enabled. It changes the default dashboard data provider ID and domain name for authorization policies from itm.<Hub monitoring server name> to itm.<domain override value>. The value may not exceed 124 characters. You should configure a domain override value for these scenarios:
 - The Hot Standby high availability feature is being used for the Hub monitoring server. By configuring a domain override value, the dashboard data provider ID and domain name will not change when the portal server is configured to connect to the new acting Hub monitoring server. If you do not configure a domain override value in this scenario, you must reconfigure the connection between the IBM Dashboard Application Services Hub and the dashboard data provider and update any domain specific authorization policies when the portal server is configured to connect to the new acting Hub monitoring server.

- You have multiple Hub monitoring servers that are using a common set of authorization policies for controlling dashboard access and you want to create some domain specific authorization policies. You should specify a domain override value for this scenario if you want to use a more user friendly domain name in your authorization policies than the default value of itm.<Hub monitoring server name>.
- If you modify the domain override after you have configured a connection in your Dashboard Application Services Hub Server to the dashboard data provider then you must delete the connection and re-add it. See the *IBM Tivoli Monitoring Administrator's Guide* for details on how to configure a dashboard data provider connection. Also if you have created any domain specific authorization policies using the default domain name then you must delete the permissions that use the previous domain name and create new permissions that use the new domain name when you change the domain override value. See the *IBM Tivoli Monitoring Command Reference* for details on the tivcmd CLI commands used to create and work with authorization policies.

Configure Tivoli Enterprise Portal Server

TEMS Connection | Agent Parameters | **Dashboard data provider**

☒ Enable the dashboard data provider

Domain override

☒ Enable authorization policies

Authorization Policy Server connection information

Hostname or IP address

Protocol

☒ HTTPS

☐ HTTP

Port

Polling interval (minutes)

Policy store expiration interval days hours

User ID

Password

Confirm password

☐ Edit host specific configuration

Host

Save Reload Help Cancel

Figure 49. Configure Dashboard data provider

10. Select the **Enable authorization policies** check box if you want to use authorization policies to control which managed systems and managed system groups a user can access in monitoring

dashboards of the IBM Dashboard Application Services Hub. Once authorization policies are enabled, only authorized users will see data in monitoring dashboards. You should only enable authorization policies if your administrators have already created the initial set of policies for dashboard user access using the Tivoli Authorization Policy Server and tivmd Command Line Interface for Authorization Policy. See the *Enabling role-based authorization policies* topic in the *IBM Tivoli Monitoring Administrator's Guide* for details on how to create and work with authorization policies.

Note: If authorization policies have not been created yet, leave this box unchecked. After the authorization policies are ready reconfigure the portal server and enable authorization policies. The fields on the configuration window are described in the following table:

Table 47. Configuration information for the Authorization Policy Server

Field	Description
Hostname or IP Address	IP Address or fully qualified hostname of the IBM Dashboard Application Services Hub with the Authorization Policy Server. This parameter is required.
Protocol	Choose the protocol used to connect to the IBM Dashboard Application Services Hub with the Authorization Policy Server. The default value is HTTPS. This parameter is not required.
Port	Choose the port used to connect to the IBM Dashboard Application Services Hub with the Authorization Policy Server. The default value is 16311 for the HTTPS protocol and 16310 for the HTTP protocol. The valid port values are from 1 to 65535 inclusive. This parameter is not required.
Polling Interval	How often the local data store is updated from the Authorization Policy Server by the policy client running on the portal server. The default is 30 minutes. Valid values are from 5 to 1440 minutes inclusive. This parameter is not required.
Policy Store Expiration Interval	If the policy store cannot be updated from the Authorization Policy Server, this interval is the amount of time the local policy store will continue to be utilized from the last update. If the Authorization Policy Server cannot be accessed for the time interval specified by this parameter, all subsequent requests for dashboard data will fail with an authorization error until the Authorization Policy Server is available again. The default is 7 days and 0 hours. The value specified for hours must be in the range of 0-23 hours. If the expiration interval is set to 0 days and 0 hours, the policy store will never expire. This parameter is not required.
User ID	The name of the user that the portal server will use to access the IBM Dashboard Application Services Hub with Authorization Policy Server. This user must be added to the PolicyDistributor authorization policy core role or to a custom role that has been granted permission to perform the distribute operation for the role object type. See the <i>Enabling role-based authorization policies</i> chapter in the <i>IBM Tivoli Monitoring Administrator's Guide</i> and the <i>IBM Tivoli Monitoring Command Reference</i> for more details on creating authorization policy roles and assigning users to them. This parameter is required.

Table 47. Configuration information for the Authorization Policy Server (continued)

Field	Description
Password	The password for the user. This parameter is required.
Confirm Password	Confirm the password by entering it again. This parameter is required.

Enter the required information for the Authorization Policy Server in the fields provided and click **OK**. For more information about the Tivoli Authorization Policy Server and Tivoli Authorization Policy Command-Line Interface, see “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318.

11. When you are done configuring the portal server, click **Save** to save your settings and close the window.
12. If you selected **Validate User with LDAP** in step 5.c, enter the LDAP configuration details in the fields provided and click **OK**. For more information about these parameters, see *Prerequisites for configuring LDAP authentication on the portal server* in the *IBM Tivoli Monitoring Administrator's Guide*.

Figure 50. LDAP Configuration

13. If you selected **Enable Single Sign On** in step 5.d the Single Sign On window is displayed. Enter values for the following fields and click **OK**:

- Domain name
- Realm name

You can also export or import LTPA keys from other participating single sign-on applications. For more information about configuring these fields and setting up SSO, see *Prerequisites for configuring LDAP authentication on the portal server* and *About single sign-on* in the *IBM Tivoli Monitoring Administrator's Guide*.

Starting the portal server

From the bin directory of /opt/IBM/ITM (or where you installed IBM Tivoli Monitoring), run the following command to start the portal server:

```
./itmcmd agent start cq
```

Upgrading a 32 bit portal server to 64 bit

The installation process does not convert an existing 32-bit Tivoli Enterprise Portal Server running on Linux or UNIX to 64 bit (although it does convert an existing 32-bit portal server from a prior release to a 32-bit portal server for the current release).

About this task

The following procedure allows you to manually perform this conversion from 32 bit to 64 bit.

Procedure

1. Using the product media for the current release, install and configure the 32-bit Tivoli Enterprise Portal Server for IBM Tivoli Monitoring.
Before proceeding to step 2, ensure that the portal server is either completely initialized or completely stopped.
2. Invoke `tepsBackup` from the `/ITM_installdir/bin` directory:

```
cd /ITM_installdir/bin
./tepsBackup
```


where:
ITM_installdir
is the root location of your IBM Tivoli Monitoring environment.
This process creates a compressed tar file with default and customized application data in the `/ITM_installdir/tmp` directory.
3. Uninstall the 32-bit portal server. If asked to completely remove the Tivoli Monitoring installation directories, answer **no**; otherwise the backup tar file will be lost.
4. Again using the product media for the current release, install and configure the 64-bit Tivoli Enterprise Portal Server.
Before proceeding to step 5, ensure that the portal server is either completely initialized or completely stopped.
5. Invoke `tepsRestore` from the `/ITM_installdir/bin` directory:

```
cd /ITM_installdir/bin
./tepsRestore
```

Installing monitoring agents

This section describes how to install distributed monitoring agents. A *distributed* monitoring agent is one that is installed on a distributed (not z/OS) operating system. These instructions also apply to the Warehouse Proxy and Summarization and Pruning Agents.

For instructions on how to install a z/OS monitoring agent, see the agent documentation for your z/OS agent product.

To install a distributed monitoring agent, use the appropriate installation media:

- Use the *IBM Tivoli Monitoring Agent* DVD or DVD image to install the agents in the following list
 - Linux OS
 - UNIX OS
 - Windows OS
 - Agentless OS
- You can also use the *IBM Tivoli Monitoring Base* DVD to install application support for any IBM Tivoli Monitoring V6.x-based distributed monitoring agent you intend to install. Installing and configuring

application support on base components (monitoring servers, portal server, desktop clients) when you install them means that you do not have to stop and restart those base components later when you install the agent and its application support.

- Use the agent product DVDs to install distributed monitoring agents that are delivered separately from the IBM Tivoli Monitoring base installation package. For example, use the IBM Tivoli Monitoring for Databases product DVDs to install a monitoring agent for DB2 Database for Linux, UNIX, and Windows or for Oracle. Depending on the agent that you are installing, there might be additional configuration steps required. See the agent documentation for more information.
- There are a set of best practice situations for the version 6.2.3 or later UNIX OS agent, Linux OS agent, and Windows OS agent. These situations are automatically installed, but not automatically distributed, during installation to the default managed system group or managed system. The best practice situations can be recognized by the `_BP_` string in the situation name.

Note: If you upgrade any OS agent to V6.3 you must also separately upgrade the Tivoli Enterprise Services User Interface Extensions (tacmd CLI) on the same machine to V6.3. If tacmd CLI (component code ue) is not upgraded, commands run from this machine will fail.

All monitoring agents require that agent-specific application support be installed on the monitoring servers, portal server, and portal desktop clients. See “Installing and enabling application support” on page 275 for information.

The following sections provide instructions for installing a monitoring agent:

- “Windows: Installing a monitoring agent”
- “Linux or UNIX: Installing a monitoring agent” on page 267

Windows: Installing a monitoring agent

Use the following steps to install a distributed monitoring agent on a Windows computer.

Procedure

1. If you are installing on either an x86-32 CPU or an x86-64 CPU, launch the installation wizard by double-clicking the `setup.exe` file in the `\WINDOWS` subdirectory on the installation media. However, if you are installing on an Itanium CPU, launch the wizard by double-clicking on the `setup.exe` file in the `\WIA64` directory.

Use either the *IBM Tivoli Monitoring Agent* DVD or a distributed agent product CD for Windows. Do *not* use a Data Files for z/OS CD.

2. Click **Next** on the Welcome window.

Note: If you have another IBM Tivoli Monitoring component already installed on this computer, select **Modify** on the Welcome window to indicate that you are updating an existing installation. Click **OK** on the message telling you about preselected items. Then skip to Step 6 on page 263.

3. Read and accept the software license agreement by clicking **Accept**.
4. Choose the directory where you want to install the product. Click **Next**.

Note:

- a. This step applies only to those agents that you install from the IBM Tivoli Monitoring installation image. Agents installed from the agent installation image do not have this step. If you are installing an agent from an agent installation image, skip to step 6 on page 263.
- b. If you specify an incorrect directory name, you will receive the following error:

The IBM Tivoli Monitoring installation directory cannot exceed 80 characters or contain non-ASCII, special or double-byte characters.

The directory name can contain only these characters:

"abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ _\:0123456789()~-./".

5. Type a 32-character encryption key. This key must be the same as the key that was used during the installation of the monitoring server to which this monitoring agent connects.
Click **Next** and then click **OK** to confirm the encryption key.
6. On the Select Features window, expand **Tivoli Enterprise Monitoring Agents**.
7. Select the names of the agents that you want to install and click **Next**.

Note:

- a. Agents available on the IBM Tivoli Monitoring installation image include the Warehouse Proxy Agent and the Summarization and Pruning Agent. Before you install these agents, follow the instructions later in this book for setting up a Tivoli Data Warehouse solution, beginning with Chapter 19, "Tivoli Data Warehouse solutions," on page 509.
- b. If both of these conditions are met:
 - The components you selected for installation would result in a mixed 32-bit and 64-bit environment.
 - No Agent Compatibility Package (see "Installing the Agent Compatibility (AC) component" on page 265) can be found on the installation DVD, or the package found is not at the appropriate level.

then this error message is displayed:

The Agent Compatibility Package in version *ver_num* is required but unavailable.

8. If a monitoring server is not installed on this computer, go to Step 9.

If you are installing monitoring agents on a computer that already has a monitoring server installed, the Agent Deployment window is displayed.

The Agent Deployment window lists monitoring agents on this installation image that you can add to the agent depot. The agent depot contains agents that you can deploy to remote computers. For information about how to deploy agents in the agent depot to remote computers, see Chapter 11, "Deploying monitoring agents across your environment," on page 371.

Note: By default, the agent depot is located in the *itm_installdir\CMS\depot* directory on Windows. If you want to use a different directory, create the directory (if it does not exist), and specify the directory using the DEPOTHOME key in the KBBENV file.

Select the agents, if any, that you want to add to the agent depot. (You can add agents to the agent depot at a later time by updating your installation.) IBM strongly recommends you also add the agent compatibility (AC) package at this time. Click **Next**.

9. This step applies only to agents that you install from the IBM Tivoli Monitoring installation image. If you are installing agents from an agent product installation image, go to Step 10.

If no IBM Tivoli Monitoring component has been previously installed on this computer, a window is displayed for you to select a program folder for the Windows Start menu. Select a program folder and click **Next**. The default program folder name is IBM Tivoli Monitoring.

10. Review the installation summary details. The summary identifies what you are installing and where you chose to install it. Click **Next** to start the installation.

After installation is complete, a configuration window (called the Setup Type window) is displayed.

11. Clear the check boxes for any components that are already installed and configured (at the current release level) on this computer, unless you want to modify the configuration. (For example, clear the check box for the Tivoli Enterprise Monitoring Server if it has already been installed and configured on this computer.) Click **Next** to start configuring all selected components.
12. Define the communications between the monitoring agents and the monitoring server:
 - a. If the agents must cross a firewall to access the monitoring server, select **Connection must pass through firewall**.
 - b. Identify the type of protocol that the agents use to communicate with the monitoring server.

You have four choices: IP.UDP, IP.PIPE, IP.SPIPE, or SNA. You can specify up to three methods for communication. If the method you have identified as Protocol 1 fails, Protocol 2 is used. If Protocol 2 fails, Protocol 3 is used. When you configure the protocol for the Warehouse Proxy Agent, configure the same protocols used by the application agents and by the hub monitoring. If the proxy agent does not have the same protocol as the hub monitoring server, it cannot register with the hub. If the proxy does not have the same protocol as the application agents, then the application agents cannot communicate with the proxy when they to create a route to it.

Note: Do not select **Optional Secondary TEMS Connection**. You can set up the failover support for agents after installation, as described in the *IBM Tivoli Monitoring High Availability Guide for Distributed Systems*.

- c. Click **OK**. A second configuration window is displayed.
- d. Complete the fields shown in the following tableTable 48 for the protocol that you specified.

Table 48. Communications protocol settings

Field	Description
IP.UDP Settings	
Hostname or IP Address	The host name or IP address for the hub monitoring server.
Port # or Port Pools	The listening port for the hub monitoring server. The default number is 1918.
IP.PIPE Settings	
Hostname or IP Address	The host name or IP address for the hub monitoring server.
Port Number	The listening port for the monitoring server. The default number is 1918.
IP.SPIPE Settings	
Hostname or IP Address	The host name or IP address for the hub monitoring server.
Port number	The listening port for the hub monitoring server. The default value is 3660.
SNA Settings	
Network Name	The SNA network identifier for your location.
LU Name	The LU name for the monitoring server. This LU name corresponds to the Local LU Alias in your SNA communications software.
LU 6.2 LOGMODE	The name of the LU6.2 LOGMODE. The default value is "CANCTDCS."
TP Name	The transaction program name for the monitoring server.
Local LU Alias	The LU alias.

- e. Click **OK** to exit the Configuration Defaults for Connecting to a TEMS window.

For additional information about these parameters, press the **Help** button.

13. Click **Finish** to complete the installation.
14. Click **Finish** on the Maintenance Complete window if you are updating an existing installation.
15. Open the Manage Tivoli Enterprise Monitoring Services utility (if it does not open automatically) to see if the monitoring agents that you installed have been configured and started. If **Yes** is displayed in the **Configured** column, the agent has been configured and started during the installation process.
16. If the value in the **Configured** column is blank and **Template** is displayed in the **Task/Subsystem** column, right-click the Template agent and complete the following steps:

- a. Click **Configure Using Defaults**.
- b. Complete any windows requiring information by using the agent-specific configuration settings in the user's guide for your agent

Note: Do not type non-ASCII characters on any of these windows. Typing characters from other character sets has unpredictable results.

- c. Repeat this step as necessary to create monitoring agent instances for each application instance you want to monitor.

Installing the Agent Compatibility (AC) component

To provide a runtime environment capable of supporting mixed 32- and 64-bit executables, as well as to provide backward compatibility between native 64-bit agents and existing 32-bit agents, the 32/64 Bit Agent Compatibility Package (component code AC) is used.

The AC package comprises the following pieces:

- 32-bit and 64-bit agent frameworks
- 32-bit and 64-bit versions of the IBM Global Security Toolkit (GSKit), component KGS
- The Tivoli Monitoring configuration utilities, component KGL

Note: The AC component is a Windows component only. There is no equivalent for Linux or UNIX systems.

When to install the AC component

You need to install the AC component in these situations:

- When adding a native 64-bit agent into preexisting IBM Tivoli Monitoring installations comprising only 32-bit components.
- When adding a 32-bit agent into preexisting Tivoli Monitoring installations comprising only 64-bit components.
- The Tivoli Monitoring installer has detected a situation where proceeding without the AC component would result in incompatibility either between the 32-bit and the 64-bit agent frameworks or between the 32-bit and the 64-bit GSKit libraries.

The IBM Tivoli Monitoring installer will automatically detect any of the above situations and then stop the installation with this error:

The version *ver_num* or newer of the Agent Compatibility Package must be installed in order to ensure the correct cooperation between 32bit and 64bit components. Exiting installation.

When not to install the AC component

The AC component can be installed with any 32-bit or 64-bit agent where a mixed 32/64 bit environment is anticipated. There is no need to install the AC component if only 32-bit or only 64-bit IBM Tivoli Monitoring agents are planned for this machine.

Installing the AC component using the Windows GUI

The AC component can be found on the Agents DVD. IBM recommends you install the AC component at the same version as the Windows OS agent (component code NT). Since the AC is a standard Tivoli Monitoring component, it can be installed using the standard interactive IBM Tivoli Monitoring installer by selecting the 32/64 Bit Agent Compatibility Package feature for installation, as shown in Figure 51 on page 266.




Figure 51. Installing the Agent Compatibility Package (component code AC)

Remotely deploying the AC components

There are special considerations for deploying native 64-bit agents as well as the AC component. When deploying an agent to a Windows machine running an x86-64 CPU, checks are automatically made to verify whether the AC component is required. If the checks report the AC component is needed and it is available in the depot, it is sent automatically with no action required on your part. However, if the checks report the AC component is needed but the component is *not* available in the depot, an error is reported, and the deployment request fails. Therefore, it is highly recommended that you populate the deployment depot with the AC component.

Sample `tacmd addBundles` command to add the AC component to the deploy depot:

```
tacmd addBundles -i C:\ITM_Agents_Image\WINDOWS\Deploy -t ac
```

For more details on managing the agent depot, see Chapter 11, “Deploying monitoring agents across your environment,” on page 371.

Note: Once you have added the AC bundle to the remote-deployment depot, it is listed among the available packages in the Tivoli Enterprise Portal. Thus your interactive users can select it when invoking the `tacmd AddSystem` command for a particular 64-bit node. However, if they do this, your users will receive this error:

```
KFWITM291E an Agent Configuration schema not found
```

When it becomes necessary to invoke the remote deployment of the Agent Compatibility package, instead use the CLI.

Installing the Embedded Java Runtime and the User Interface Extensions

On nodes where only the Windows OS agent has been installed, either locally or remotely, the Embedded Java Runtime and the Tivoli Enterprise Services User Interface Extensions (the KUE component) are not also installed by default.

If you later decide to install an Agent Builder agent on this node, you may not be able to reconfigure this factory agent, and you may receive the error shown in Figure 52.




Figure 52. Java Runtime Environment Not Detected error

This error also may occur when you attempt to run a `tacmd` CLI command on nodes where either the Embedded Java Runtime or the User Interface Extensions are unavailable.

If this happens, you must install the Embedded Java Runtime and the KUE component. Complete one of the following procedures, depending on the installation method you are following:

local GUI installation

Select Tivoli Enterprise Services User Interface Extensions from the list of features to install.

local silent installation

Uncomment this line in the silent response file:

```
;KUEWICMA=Tivoli Enterprise Services User Interface Extensions
```

remote installation

First ensure that component UE has been added to the server depot. Once the UE component is available, from the command-line, perform a **`tacmd addsystem`** command, specifying `-t UE` as the component to install.

Once you have completed one of the above procedures, the Embedded Java Runtime and the Tivoli Enterprise Services User Interface Extensions are installed and can be accessed on this node.

Linux or UNIX: Installing a monitoring agent

Use the following steps to install and configure a distributed monitoring agent on a Linux or UNIX computer.

Table 49. Steps for installing a monitoring agent on Linux or UNIX

Steps	Where to find information
Install the monitoring agent.	"Installing the monitoring agent" on page 268
Configure the monitoring agent.	"Configuring the monitoring agent" on page 269 Some agents require additional, agent-specific configuration parameters. See the agent documentation for the specific agents that you are configuring.
Change the file permissions for files on the computer where you installed the agent.	"Postinstallation steps for nonroot installations" on page 270

Table 49. Steps for installing a monitoring agent on Linux or UNIX (continued)

Steps	Where to find information
Start the monitoring agent.	"Starting the monitoring agents" on page 271

Installing the monitoring agent

Use the following steps to install a monitoring agent on a Linux or UNIX computer.

Procedure

1. In the directory where you extracted the installation files, run the following command:
./install.sh
2. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory (/opt/IBM/ITM) or type the full path to a different directory.

Note: You must not specify the path of the directory containing ./install.sh as your IBM Tivoli Monitoring home directory. On certain platforms, this can cause the plug-in JAR files to overwrite themselves and become zero length files. The installation will fail as a result.

In addition, do not use the number sign (#) or exclamation mark (!) characters anywhere in your installation directory path.

3. If the installation directory does not already exist, you are asked if you want to create it. Type y to create this directory and press Enter.
4. The following prompt is displayed:
Select one of the following:
1) Install products to the local host.
2) Install products to depot for remote deployment (requires TEMS).
3) Install TEMS support for remote seeding
4) Exit install.

Please enter a valid number:

Note: This prompt might vary depending on the installation image from which you are installing. Type 1 to start the installation and display the software license agreement.

5. Press Enter to read through the agreement.
6. Type 1 to accept the agreement and press Enter.
7. Type a 32 character encryption key and press Enter. This key should be the same as the key that was used during the installation of the monitoring server to which this monitoring agent connects.

Note: This step applies only to those agents that you install from the IBM Tivoli Monitoring installation image. Agents installed from the agent installation image do not need to provide the encryption key.

A numbered list of available operating systems is displayed.

8. Type 1 to install the IBM Tivoli Monitoring support for your current operating system. Press Enter. A numbered list of available components is displayed.
9. Type the number that corresponds to the monitoring agent or agents that you want to install. If you want to install more than one agent, use a comma (,) or a space to separate the numbers for each agent. Press Enter.

Note: Before you install the Warehouse Proxy Agent or Summarization and Pruning Agent, follow the instructions later in this book for setting up a Tivoli Data Warehouse solution, beginning with Chapter 19, "Tivoli Data Warehouse solutions," on page 509.

A list of the components to be installed is displayed.

10. Type 1 to confirm the installation.

The installation begins.

11. After all of the components are installed, you are asked whether you want to install additional products or product support packages. Type 2 and press Enter.
12. If your IBM Tivoli Monitoring environment is not already secured you will be asked at this point if you want to secure it. If your IBM Tivoli Monitoring environment is already secured this question is skipped. The product installation process creates the majority of directories and files with world write permissions. IBM Tivoli Monitoring provides the secureMain utility to help you keep the monitoring environment secure. You can secure your installation now, or manually execute the secureMain utility later. For more information, see Appendix G, "Securing your IBM Tivoli Monitoring installation on Linux or UNIX," on page 925.

What to do next

Continue with "Configuring the monitoring agent."

Configuring the monitoring agent

Use the following steps to configure your monitoring agent.

Procedure

1. Run the following command:

```
./itmcmd config -A pc
```

where *pc* is the product code for your agent. For the UNIX agent, use the product code "ux"; for Linux, use "lz". See Appendix D, "IBM Tivoli product, platform, and component codes," on page 883 for a list of agent product codes.

2. Press Enter when you are asked if the agent connects to a monitoring server.
3. Type the host name for the monitoring server.
4. Type the protocol that you want to use to communicate with the monitoring server. You have four choices: ip.udp, sna, ip.spipe, or ip.pipe. Press Enter to accept the default protocol (IP.PIPE).
5. If you want to set up a backup protocol, enter that protocol and press Enter. If you do not want to use backup protocol, press Enter without specifying a protocol.
6. Depending on the type of protocol you specified, provide the following information when prompted:

Table 50. UNIX monitoring server protocols and values

Protocol	Value	Definition
IP.UDP	IP Port Number	The port number for the monitoring server. The default is 1918.
SNA	Net Name	The SNA network identifier for your location.
	LU Name	The LU name for the monitoring server. This LU name corresponds to the Local LU Alias in your SNA communications software.
	Log Mode	The name of the LU6.2 LOGMODE. The default value is "CANCTDCS."
IP.PIPE	IP.PIPE Port Number	The port number for the monitoring server. The default is 1918.
IP.SPIPE	IP.SPIPE Port Number	The port number for the monitoring server. The default is 3660.

7. Press Enter to *not* specify the name of the KDC_PARTITION.

8. Press Enter when you are asked if you want to configure the connection to a secondary monitoring server. The default value is no.
9. Press Enter to accept the default for the Optional Primary Network Name (none).

What to do next

You must complete the configuration of the Warehouse Proxy Agent using the Manage Tivoli Enterprise Monitoring Services graphical user interface (GUI), which requires an X11 GUI interface. See the instructions for configuring the Warehouse Proxy for the appropriate Data Warehouse:

- DB2 for Linux, UNIX, and Windows: “Configuring a Warehouse Proxy Agent on Linux or UNIX (JDBC connection)” on page 555
- SQL Server: “Configuring a Warehouse Proxy Agent on Linux or UNIX (JDBC connection)” on page 605
- Oracle: “Configuring a Warehouse Proxy Agent on Linux or UNIX (JDBC connection)” on page 629

If you used a non-root user to install a monitoring agent on a UNIX computer, the file permissions are initially set to a low level. Use the procedure in “Postinstallation steps for nonroot installations” to change these file permissions:

Postinstallation steps for nonroot installations

If you used a non-root user to install a monitoring agent on a UNIX computer, complete the following procedure.

Procedure

1. Log in to the computer as the root user.
2. Create a new group (such as *itmusers*) to own all of the files in the IBM Tivoli Monitoring installation directory. For an AIX computer, use the `mkgroup` command. For Linux, Solaris, and HP-UX computers, use the `groupadd` command. See the documentation for your operating system for the complete syntax information.
3. Add both the root and the nonroot user to the new group *itmusers*. For an AIX computer, use the `chgrpmm` command. For Linux, Solaris, and HP-UX computers, use the `usermod` command. See the documentation for your operating system for the complete syntax information.
4. Run the following commands and verify that both the root and the nonroot user are in the *itmusers* group:

```
groups root
groups nonroot
```
5. Run the following command to set the permissions and group ownership of CANDLEHOME:

```
CANDLEHOME/bin/secureMain -g itmusers lock
```
6. Run the following command to update the system boot scripts:

```
CANDLEHOME/bin/UpdateAutoRun.sh
```
7. Start the agent as described in “Starting the monitoring agents” on page 271.

What to do next

- Whenever you update the agent, you must log in to the system as the root user and run the `secureMain` step again. If the agent is running:
 1. Stop the agent.
 2. Run `secureMain` as described in Step 5.
 3. Restart the agent.
- If you installed IBM Tivoli Monitoring using a nonroot user and did not run `secureMain` as described in Step 5, you must run `secureMain` before using the **`tacmd updateagent`** command to update the agent. If you do not run `secureMain` first, the update process will not complete.

- On AIX systems you should invoke the AIX `slibclean` command:

```
su -c "/usr/sbin/slibclean"
```

Starting the monitoring agents

You can either start all agents running on a computer or start individual agents by using the product codes.

To start all monitoring agents, run the following command:

```
./itmcmd agent start all
```

To start specific agents, run the following command:

```
./itmcmd agent start pc pc pc
```

where *pc* is the product code for the agent that you want to start. See Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883 for a list of agent product codes.

To start multi-instance agents (agents that may run more than one instance on a computer, like Seibel or Domino® agents), run the following command:

```
./itmcmd agent -o instance_name start pc
```

where *pc* is the product code for the agent that you want to start and *instance_name* is the name that uniquely identifies the instance you want to start. See Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883 for a list of agent product codes.

Populating the data warehouse's ManagedSystem table

If your site runs the Tivoli Data Warehouse, each time you install one or more monitoring agents, you must update the warehouse's ManagedSystem table; the `populate_agents.sql` script is provided for this purpose.

- If your site uses DB2 Database for Linux, UNIX, and Windows to manage its Tivoli Data Warehouse, call this stored procedure:

```
db2 call ITMUSER.POPULATE_OSAGENTS();
```
- If your site uses Microsoft SQL Server to manage its Tivoli Data Warehouse, run this script at the MS SQL command-line:

```
sqlcmd -i populate_agents.sql [-U my_username -P my_password] [-H my_host])
```
- If your site uses Oracle to manage its Tivoli Data Warehouse, start this procedure:

```
POPULATE_OSAGENTS('ITMUSER');
```

where:

ITMUSER

Is the database user for Tivoli Data Warehouse.

my_username

Is the user ID for the Tivoli Data Warehouse (in most cases the same as ITMUSER).

my_password

Is the password for the user.

my_host

Is the optional computer name.

Installing the Tivoli Enterprise Portal desktop client

There are two methods of deploying the desktop client: You can install the desktop client from the installation media and run and maintain it on the local system. You can also use IBM Web Start for Java to download and run the desktop client from the Tivoli Enterprise Portal Server.

See “Tivoli Enterprise Portal client” on page 27 for a discussion of the advantages and disadvantages of each type of portal server client available with IBM Tivoli Monitoring.

This section describes how to install the desktop client from the installation media on Windows and Linux computers. When you install the desktop client from the installation media, you must also install support for all the applications that you will be using.

- “Windows: Installing the desktop client”
- “Linux: Installing the desktop client” on page 273

See “Using Web Start to download and run the desktop client” on page 367 for instructions on downloading and running the desktop client from the portal server.

Windows: Installing the desktop client

Complete the following steps to install the Tivoli Enterprise Portal desktop client from the Base Infrastructure DVD or DVD image.

Procedure

1. On the computer where you want to install the desktop client, start the installation wizard by launching the setup.exe file in the \WINDOWS subdirectory on the installation media.
2. Click **Next** on the Welcome window.

Note: If you have another IBM Tivoli Monitoring component already installed on this computer, select **Modify** on the Welcome window to indicate that you are updating an existing installation. Click **OK** on the message telling you about preselected items. Then skip to Step 6.
3. Read and accept the software license agreement by clicking **Accept**.
4. Specify the directory where you want to install the portal desktop client software and accompanying files. The default location is C:\IBM\ITM. Click **Next**.

Note: If you specify an incorrect directory name, you will receive the following error:

The IBM Tivoli Monitoring installation directory cannot exceed 80 characters or contain non-ASCII, special or double-byte characters.

The directory name can contain only these characters:

"abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ _\:0123456789()~-. /".

5. Type an encryption key to use. This key must be the same as what was used during the installation of the hub monitoring server to which the client will connect. Click **Next** and then **OK** to confirm the encryption key.
6. On the Select Features window, select **Tivoli Enterprise Portal Desktop Client** from the list of components to install.

When you select the **Tivoli Enterprise Portal Desktop Client** check box, all of the check boxes in the attached subtree are automatically selected. The support check boxes in the subtree are for installing application support files for distributed monitoring agents to the portal desktop client. It is best to leave all of the support check boxes selected so you do not need to reconfigure application support as new agent types are added to your environment. However, if you purchased an OMEGAMON XE monitoring product and you did not separately purchase IBM Tivoli Monitoring, you should not install support for the distributed operating system monitoring agents. If you are installing from the Base DVD, you may see application support for components that are not supported on the operating system. For detailed information about application support, see “Installing and enabling application support” on page 275.

Note: If you are updating an existing installation (you selected **Modify** on the Welcome window), all check boxes on the Select Features window reflect your choices during the initial installation. Clearing a check box has the effect of *uninstalling* the component. Clear a check box only if you want to remove a component.

7. If you want to view IBM Tivoli Enterprise Console events through the Tivoli Enterprise Portal, expand **Tivoli Enterprise Portal Desktop Client** and ensure that **TEC GUI Integration** is selected.

8. Click **Next**.

9. If a monitoring server is not installed on this computer, go to Step 10.

If you are installing the desktop client on a computer that already has a monitoring server installed, the Agent Deployment window is displayed.

The Agent Deployment window lists monitoring agents on this installation image that you can add to the agent depot. The agent depot contains agents that you can deploy to remote computers. For information about how to deploy agents in the agent depot to remote computers, see Chapter 11, “Deploying monitoring agents across your environment,” on page 371.

Note: By default, the agent depot is located in the *itm_installdir*\CMS\depot directory on Windows. If you want to use a different directory, create the directory (if it does not exist) and specify the directory using the DEPOTHOME key in the KBBENV file.

Select the agents, if any, that you want to add to the agent depot. (You can add agents to the agent depot at a later time by updating your installation.) Click **Next**.

10. If no IBM Tivoli Monitoring component has been previously installed on this computer, a window is displayed for you to select a program folder for the Windows Start menu. Select a program folder and click **Next**. The default program folder name is IBM Tivoli Monitoring.

11. Review the installation summary details. The summary identifies what you are installing and where you chose to install it. Click **Next** to start the installation.

After installation is complete, a configuration window (called the Setup Type window) is displayed.

12. Clear the check boxes for any components that are already installed and configured (at the current release level) on this computer, unless you want to modify the configuration. (For example, clear the check box for the Tivoli Enterprise Monitoring Server if it has already been installed and configured on this computer.) If the desktop client is being installed on the same host as the portal server but as part of a separate installation process, you must reconfigure the portal server.

Click **Next** to start configuring all selected components.

13. Type the host name of the portal server and click **OK**.

14. Click **Finish** to complete the installation.

Linux: Installing the desktop client

Use the following steps to install the Tivoli Enterprise Portal desktop client.

Procedure

1. In the directory where you extracted the installation files, run the following command:
`./install.sh`
2. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory (/opt/IBM/ITM) or type the full path to a different directory.

Note: You must not specify the path of the directory containing `./install.sh` as your IBM Tivoli Monitoring home directory. On certain platforms, this can cause the plug-in JAR files to overwrite themselves and become zero length files. The installation will fail as a result.

In addition, do not use the number sign (#) or exclamation mark (!) characters anywhere in your installation directory path.

3. If the installation directory does not already exist, you are asked if you want to create it. Type 1 to create this directory and press Enter.
4. The following prompt is displayed:
 Select one of the following:
 1) Install products to the local host.
 2) Install products to depot for remote deployment (requires TEMS).
 3) Install TEMS support for remote seeding
 4) Exit install.

 Please enter a valid number:
 Type 1 to start the installation and display the software license agreement.
5. Press Enter to read through the agreement.
6. Type 1 to accept the agreement and press Enter.
7. Type an 32-character encryption key to use and press Enter. This key should be the same key as that used during the installation of the portal server to which the client will connect.
 A numbered list of available operating systems is displayed.
8. Type 3 to install the desktop client support for your current operating system. Press Enter.
 A message is displayed indicating that the Tivoli Enterprise Portal Desktop Client is about to be installed.
9. Type 1 to confirm the installation.
 The installation begins.
10. After the portal desktop client is installed, you are asked whether you want to install additional products or product support packages. Type 1 and press Enter.
 A numbered list is displayed, including the following application support package:
 Tivoli Enterprise Portal Desktop Client support
11. Install the application support package for the portal desktop client.
 All monitoring agents require that application support files be installed on the monitoring servers (hub and remote), portal server, and portal desktop clients in your environment. Application support files contain the information required for agent-specific workspaces, helps, predefined situations, and other data.

 This step installs the application support files for base monitoring agents. The *base* monitoring agents are included with the base IBM Tivoli Monitoring installation package. For detailed information about application support, see "Installing and enabling application support" on page 275.
 - a. Type the number that corresponds to Tivoli Enterprise Portal Desktop Client support and press Enter.
 A numbered list of base monitoring agents is displayed.
 - b. Type the numbers of the base monitoring agents for which you want to install application support, or type the number that corresponds to All of the above. Type the numbers on the same line separated by spaces or commas. Press Enter.

 It is best to select all of the base monitoring agents (All of the above) so you do not need to reconfigure application support as new agent types are added to your environment.
 - c. Type 1 to confirm the installation and press Enter.
 The installation begins.

Note: This step installs the application support files. However, you must *enable* the application support by configuring the portal desktop client. The next sections shows you how to configure the portal desktop client.
12. After application support for the monitoring agents is installed, you are asked whether you want to install additional products or product support packages. Type 2 and press Enter.

What to do next

The next step is to configure the desktop client. Use the instructions in “Linux: Configuring the desktop client.”

Linux: Configuring the desktop client

Complete the following steps to configure the desktop client if you installed the client from the IBM Tivoli Monitoring installation media. You do not need to complete this procedure if you obtained the desktop client by using IBM Web Start for Java to download it from the Tivoli Enterprise Portal Server.

Procedure

1. At the command-line `/opt/IBM/ITM/bin` directory (or the `/bin` subdirectory where you installed the product), run the following command:

```
./itmcmd config -A cj
```
2. Press Enter to use the default instance name.
3. Type the host name for the portal server and press Enter.
4. Press Enter when you are asked if you want to use HTTP Proxy support. The default value is no.
5. Start the desktop client:

```
/itmcmd agent start cj
```

Installing and enabling application support

Before you can view data collected by monitoring agents, you must install and enable application support for those agents. Application support files provide agent-specific information for workspaces, helps, situations, templates, and other data.

Application support for a monitoring agent includes two types of files:

- **SQL files** are required for adding product-provided situations, templates, and policies to the Enterprise Information Base (EIB) tables maintained by the hub monitoring server. These SQL files are also called *seed data*, and installing them on a monitoring server is also called *seeding* the monitoring server.
- **Catalog and attribute (cat and atr) files** are required for presenting workspaces, online help, and expert advice for the agent in Tivoli Enterprise Portal.

Note: There is a limit of 512 catalog and attribute files that can reside on a Tivoli Enterprise Monitoring Server. When this limit is reached, the Tivoli Enterprise Monitoring Server will not start and the message *Number of installed packages exceeds the maximum of 512* is displayed in the error log. For more information, see the *IBM Tivoli Monitoring Troubleshooting Guide*.

All monitoring agents require that application support be configured on all instances of the following infrastructure components:

- Tivoli Enterprise Monitoring Server (both hub and remote monitoring servers)
- Tivoli Enterprise Portal Server
- Tivoli Enterprise Portal desktop client, if the desktop client was installed from the installation media.
You do not need to configure application support for desktop clients downloaded from the Tivoli Enterprise Portal Server using IBM Web Start for Java.

Note: The self-describing agent capability integrates the installation of an agent with the dispersal and installation of associated product support files throughout your IBM Tivoli Monitoring infrastructure. The self-describing agent feature makes it possible for new or updated IBM Tivoli Monitoring agents to become operational after installation, without having to perform additional product support installation steps or restarting the monitoring server or portal server. For more information see “Self-describing agent

installation” on page 394. However, the following components are not able to dynamically refresh application support and require a restart after application support has been installed by the self-describing agent feature:

- Tivoli Enterprise Monitoring Automation Server: If this optional component is installed with the Hub monitoring server, the automation server must be restarted after application support has been installed on the Hub monitoring server by the self-describing agent feature.
- Dashboard data provider component of the Tivoli Enterprise Portal Server: If this optional component is enabled, the portal server must be restarted after application support has been installed on the portal server by the self-describing agent feature.

Application support for monitoring agents is installed independently of where and when the monitoring agents themselves are installed:

- Install application support for a particular type of monitoring agent on the monitoring servers, portal server, and portal desktop clients. Install agents of that type on any managed system in the environment that is compatible with the agent.
- Install application support for a type of monitoring agent *before* or *after* any monitoring agents of that type are installed. After you install application support for a particular type of monitoring agent, you can add any number of agents of that type to your environment without having to install application support again.

For example, you can install application support for a Linux OS monitoring agent (the agent *type*) to a Windows monitoring server (using the IBM Tivoli Monitoring installation media for Windows). Later, you can install any number of Linux OS monitoring agents to Linux computers in your environment (using the IBM Tivoli Monitoring installation media for Linux).

Important:

1. If you are installing a non-OS agent remotely through the Tivoli Enterprise Portal (see “Deploying non-OS agents” on page 377), application support for the agent must be installed on the Tivoli Enterprise Portal Server *before* the agent is deployed.
2. When you install application support for an IBM Tivoli Composite Application Manager agent on the V6.2.2 fix pack 2 (or subsequent) version of the Tivoli Enterprise Monitoring Server, do not select for installation the File Transfer Enablement component. If you do, the File Transfer Enablement component shipped with the V6.2.2 fix pack 2 (or subsequent) monitoring server will be replaced, and the **tacmd getfile**, **tacmd putfile**, and **tacmd executecommand** commands will become inoperative.

Because application support for some IBM Tivoli Monitoring 6.x-based distributed agents is included on the Base Infrastructure DVD, you may see application support files that do not apply to all systems.

Configuring application support is a two-step process:

1. Installing the application support files (from installation media).
2. Enabling the application support (sometimes referred to as *adding* or *activating* the application support).

On the portal server and portal desktop clients, application support is enabled when the component is configured. On monitoring servers, application support is enabled by *seeding* the database with agent-specific information.

The procedures for configuring application support differ by operating system, as summarized in Table 51 on page 277. On Windows, both installation and enablement of application support are accomplished during the installation of the monitoring servers, portal server, and desktop clients. On Linux or UNIX, this two-step process is more visible, with the enablement step done separately from the installation.

Table 51. Procedures for installing and enabling application support

Operating system	Monitoring servers	Portal server	Desktop clients ¹
Windows	Install and enable application support using installation media.	Install and enable application support using installation media.	Install and enable application support using installation media.
Linux or UNIX	<ol style="list-style-type: none">1. Install application support files from installation media.2. <i>Seed</i> the monitoring server using:<ul style="list-style-type: none">• itmcmd support² command, OR• Manage Tivoli Enterprise Monitoring Services window	<ol style="list-style-type: none">1. Install application support files from installation media.2. <i>Configure</i> the portal server using:<ul style="list-style-type: none">• itmcmd config command, OR• Manage Tivoli Enterprise Monitoring Services window	<ol style="list-style-type: none">1. Install application support files from installation media.2. <i>Configure</i> the desktop client using:<ul style="list-style-type: none">• itmcmd config command, OR• Manage Tivoli Enterprise Monitoring Services window
z/OS	<p>This book does not describe configuring application support for a monitoring server installed on a z/OS system. For more information and instructions, see <i>Configuring the Tivoli Enterprise Monitoring Server on z/OS</i> (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm).</p> <p>The portal server and desktop client are not supported on z/OS. If you have a monitoring server on z/OS, use the procedures in this book to configure application support on the portal server and desktop clients.</p>		
<ol style="list-style-type: none">1. You need to configure application support for desktop clients that are installed from the installation media. You do not need to configure application support for desktop clients that are installed by using IBM Java Web Start to download the client from the Tivoli Enterprise Portal Server.2. You can seed a nonlocal monitoring server, even if one is not installed on the local computer, by installing the support using option 3, Install TEMS support for remote seeding, then using Manage Tivoli Enterprise Monitoring Services to seed the nonlocal monitoring server. You cannot use itmcmd support to seed a nonlocal monitoring server.3. There is no way to uninstall the application support files laid down without uninstalling the Tivoli Enterprise Monitoring Server.			

Selecting the correct support media

Application support for the base monitoring agents provided with IBM Tivoli Monitoring is found on the following DVDs:

- *IBM Tivoli Monitoring Base DVD* contains the application support for the agents listed in Table 52 on page 278.
- *IBM Tivoli Monitoring Tivoli Performance Analyzer Reports and Domain Definitions DVD* contains the Performance Analyzer domain application support listed in Table 53 on page 278. This DVD is platform-nonspecific: the same DVD is used for installing on Windows, Linux, and UNIX.

To install application support for all IBM Tivoli Monitoring base monitoring and agents on distributed nodes, use the procedures documented in this chapter for installing application support on the monitoring servers, the Tivoli Enterprise Portal, and the portal desktop client. See “Configuring application support on the IBM Tivoli Monitoring Base DVDs” on page 278 for more information. For instructions on installing application support for the monitoring agents on a monitoring server on z/OS, see *Configuring the Tivoli Enterprise Monitoring Server on z/OS* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm).

Table 52. Product support on the Infrastructure and Agent DVDs

Product Code	Product Name
R3	Agentless monitor for AIX operating systems
R5	Agentless monitor for HP-UX operating systems
R4	Agentless monitor for Linux operating systems
R6	Agentless monitor for Solaris operating systems
R2	Agentless monitor for Windows operating systems
LZ	Monitoring agent for Linux OS
UL	Monitoring agent for UNIX Logs
UX	Monitoring agent for UNIX OS
NT	Monitoring agent for Windows OS
A4	Monitoring agent for IBM i
SY	Summarization and Pruning Agent
HD	Warehouse Proxy Agent
UM	IBM Tivoli Universal Agent
PA	IBM Tivoli Performance Analyzer

Table 53. Tivoli Performance Analyzer Reports and Domain Definitions

Product Code	Product Name
P0	Tivoli Performance Analyzer Domain for DB2
P3	Tivoli Performance Analyzer Domain for OS agent
P4	Tivoli Performance Analyzer Domain for Oracle
P6	Tivoli Performance Analyzer Domain for System P
PI	Tivoli Performance Analyzer Domain for ITCAM RT
PU	Tivoli Performance Analyzer Domain for VMware

For all other distributed agents, install application support from their product installation CDs. See “Configuring application support for nonbase monitoring agents” on page 279 for instructions.

Configuring application support on the IBM Tivoli Monitoring Base DVDs

You install and configure application support for the agents listed in Table 52 when you use the Base DVDs to install and configure distributed monitoring servers, the portal server, and the portal desktop clients using the instructions in this chapter.

- When you install a monitoring server on Linux or UNIX, support files for all agent types on the Base DVD are installed automatically. When you install the portal server and portal desktop client, you can specify which agent types to support.
- When you install a monitoring server, portal server, or portal desktop client on Windows, you can select the Base monitoring agents for which you want to install application support on the Select Features window. By default, all the agents are selected.

In both cases, it is best to install and enable application support for all agent types on the Base DVDs during the initial installation to avoid having to reconfigure application support on all monitoring servers, the portal server, and portal desktop clients as new agent types are added to your environment.

Note: If you install application support on the Hub monitoring server and the Tivoli Enterprise Monitoring Automation Server is also installed, you must restart the automation server after the application support has been installed.

If you need to reconfigure application support, restart the installation from the IBM Tivoli Monitoring installation media. Update the installation by installing application support for additional Base monitoring agents. When you update an installation, keep in mind the following considerations:

- Following an installation update on Linux or UNIX, you must reseed the monitoring server and reconfigure the portal server and desktop clients.
- When you update an installation on Windows, proceed through the installation as you did for the initial installation except as follows:
 - All check boxes on the Select Features window reflect your choices during the initial installation. Do not clear any check boxes. Clearing a check box has the effect of *uninstalling* the component. For example, if you are updating the installation to add application support for the IBM i agent, select only the application support check box for the IBM i agent. Do not clear any application support check boxes that are already checked unless you want to remove support for those agents.
 - You do not need to reconfigure any components (such as the monitoring server or portal server) that have already been configured.
- You must update the installation on all computers where a monitoring server, portal server, or desktop client is installed.

Before seeding actually begins, as part of the installation process, you will be asked if you want to use the seeding files that are included with the installation media or if you want to use the ones you customized during a previous installation.

Important: Do not install agent application support for the current release on a Tivoli Enterprise Monitoring Server for a prior release (for example, do not install V6.3 agent application support on a V6.2.x monitoring server). Doing so can cause agents to fail.

When the application support for an agent is updated on the Tivoli Enterprise Portal Server, the current agent DLA template file is renamed to use the .bak extension and the latest version of the template file is installed. If the agent DLA template file was edited to remove the IP address filtering for private network addresses, you must update the new version of the agent DLA template file to contain the edits that are in the .bak version of the template file. For more information on editing an agent DLA template to remove private network address filtering, see *Using the Tivoli Management Services Discovery Library Adapter* in the *IBM Tivoli Monitoring Administrator's Guide*.

Configuring application support for nonbase monitoring agents

The term *nonbase* is used here to refer to other products such as IBM Tivoli Monitoring for Applications, OMEGAMON XE monitoring agents, and IBM Tivoli Composite Application Manager monitoring agents that provide their own support files.

Some monitoring agents that run on z/OS or z/VM are packaged with their own CD that contains the data files for adding application support to distributed components. Other monitoring agents that run on z/OS are packaged with a CD that contains data files for a number of agents. If in doubt, see the configuration guide for each of your monitoring agents to find the exact name of the CD to use.

The following table shows you which installation media to use and where to find instructions for installing application support, according to the type of agent (distributed or z/OS) and whether the agent reports to a distributed or z/OS monitoring server.

Table 54. Installation media and instructions for installing application support for nonbase monitoring agents

Agent	Monitoring server	Installation media	Instructions for installing application support
Distributed	Distributed	Agent product installation CDs	Follow the instructions in this section to install application support on the monitoring server, portal server, and desktop client.
Distributed	z/OS	Agent product installation CDs	<ul style="list-style-type: none"> Follow the instructions in this section to install application support on the portal server and desktop clients. Follow the instructions in <i>Configuring the Tivoli Enterprise Monitoring Server on z/OS</i> (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm) to install application support on the monitoring servers.
z/OS	Distributed	Data Files CD	Follow the instructions in this section to install application support on the monitoring server, portal server, and desktop client.
z/OS	z/OS	Data Files CD	<ul style="list-style-type: none"> Follow the instructions in this section to install application support on the portal server and desktop clients. Follow the instructions in <i>Configuring the Tivoli Enterprise Monitoring Server on z/OS</i> (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm) to install application support on the monitoring servers.

Use the instructions in the following sections to install application support for nonbase distributed or z/OS monitoring agents on the distributed monitoring servers, portal server, and portal desktop clients in your environment:

- “Installing application support on monitoring servers”
- “Installing application support on the Tivoli Enterprise Portal Server” on page 286
- “Installing application support on the Tivoli Enterprise Portal desktop client” on page 289

Each of these sections provides information for installing application support files to a single component, such as the monitoring server. If you have multiple components on the same computer (such as the monitoring server and the portal server), combine steps from the individual sections to install application support to all components.

When the application support for an agent is updated on the Tivoli Enterprise Portal Server, the current agent DLA template file is renamed to use the .bak extension and the latest version of the template file is installed. If the agent DLA template file was edited to remove the IP address filtering for private network addresses, you must update the new version of the agent DLA template file to contain the edits that are in the .bak version of the template file. For more information on editing an agent DLA template to remove private network address filtering, see *Using the Tivoli Management Services Discovery Library Adapter* in the *IBM Tivoli Monitoring Administrator's Guide*.

Installing application support on monitoring servers

Use the following procedures to install application support for nonbase monitoring agents on distributed monitoring servers (hub or remote) in your environment.

- “Windows: Installing application support on a monitoring server” on page 281
- “Linux or UNIX: Installing application support on a monitoring server” on page 284

Windows: Installing application support on a monitoring server:

Complete the following steps to install application support for monitoring agents on Windows monitoring servers.

Before you begin

1. The monitoring server is stopped during this process.
2. If you are running in a Hot Standby environment, shut down your Hot Standby (that is, mirror) monitoring server before completing this procedure. Restart the Hot Standby monitoring server only after you have seeded the hub server.

Procedure

1. In the \WINDOWS subdirectory on the agent product CD (for distributed products) or data files CD (for z/OS products), double-click the setup.exe file to launch the installation.
2. Click **Next** on the Welcome window.

Note: If a monitoring agent is already installed on this computer, select **Modify** on the Welcome window to indicate that you are updating an existing installation. Click **OK** on the message telling you about preselected items. Then skip to Step 6.

3. On the Install Prerequisites window, read the information about the required levels of IBM Global Security Toolkit (GSKit) and IBM Java.

The check box for each prerequisite is cleared if the correct level of the software is already installed. Otherwise, the check box is selected to indicate that the software is to be installed. If you are installing support from the data files CD for z/OS agent products, you might be prompted to install Oracle Java Runtime Environment (JRE) 1.4.2, even if you have already installed IBM JRE 1.5 with the distributed components of Tivoli Management Services. The two versions can coexist, and installation of application support for some monitoring agents requires Oracle Java 1.4.2. You might also see a message indicating that you can decline the installation of JRE 1.4.2 and that accepting installation of JRE 1.4.2 results in uninstallation of other Java versions. Ignore this message, because you cannot proceed without accepting the installation of Oracle Java 1.4.2, and accepting the installation does not uninstall IBM Java 1.5.

4. Click **Next**. The prerequisite software is installed if necessary.

If the installation program installs IBM GSKit or IBM JRE, you might be prompted to restart the computer when the installation is complete. If so, you receive an abort message with a **Severe error** heading. This is normal and does not indicate a problem.

If you are prompted to reboot, do the following:

- a. Click **OK** on the window prompting you to reboot.
 - b. Click **No** on the window asking whether you want to view the abort log.
 - c. Restart the computer.
 - d. Restart the installation program.
5. Click **Accept** to accept the license agreement.
 6. If you see a message regarding installed versions being newer than the agent installation, click **OK** to ignore this message.
 7. Select the application support packages that you want to install:
 - a. On the Select Features window, select **Tivoli Enterprise Monitoring Server**.
 - b. Expand the **Tivoli Enterprise Monitoring Server** node to display a list of application support packages that you can install on the monitoring server. Initially, all application support packages are selected.
 - c. Clear the check boxes for application support packages that you do not want to install.

Note: If you are updating an existing installation (you selected **Modify** on the Welcome window), all check boxes on the Select Features window reflect your choices during the initial installation. Clearing a check box has the effect of *uninstalling* the component (for example, the Tivoli Enterprise Monitoring Server) or product package. Clear a check box for an application support package only if you want to remove the application support.

- d. If you have other components installed on the same computer, such as the desktop client, also select those components to install the component-specific application support.

e. Click **Next**.

- 8. (*Distributed agents only*) If you want to add the agent to the deployment depot, select the agent and click **Next**.

This step does not occur for z/OS agents. z/OS agents cannot be added to the deployment depot.

- 9. On the Start Copying Files window, read the list of actions to be performed and click **Next** to start the installation.

The application support packages that you selected are installed.

- 10. On the Setup Type window, do the following:
 - a. Select **Install application support for a local/remote Tivoli Enterprise Monitoring Server**.
 - b. Optionally, select the check box for launching the Manage Tivoli Enterprise Monitoring Services window. (If selected, this window is displayed when the installation procedure is finished.)
 - c. Clear the check boxes for any components that you have already installed on this computer, such as the monitoring server.
 - d. Click **Next**.

- 11. On the two Tivoli Enterprise Monitoring Server configuration windows that are displayed, make sure the information is correct and click either **Next** or **OK**.

- 12. Enable application support on the monitoring server:

In Step 7 on page 281, you selected the application support packages that you wanted to install on the monitoring server. In this step, you activate the application support through a process known as *seeding* the monitoring server.

- a. Specify the location of the monitoring server to which to add application support. You have two choices:

- **On this computer**
- **On a different computer**

Click **OK**.

For additional information about these parameters, press the **Help** button.

- b. If you are updating a hub Tivoli Enterprise Monitoring Server, you are asked to choose whether you want to add the default managed system list when you process the application-support files:

All Add the default managed system groups to all applicable situations.

New Add the default managed system groups to all applicable situations from the product support packages being seeded for the first time. Modifications are not made to managed system groups in previously upgraded product support packages.

None The default managed system group is not added to any situation.

Note: Not all situations support the default managed group setting. For some, you might need to manually define the distribution using the Tivoli Enterprise Portal due to the specific content of the agent support package.

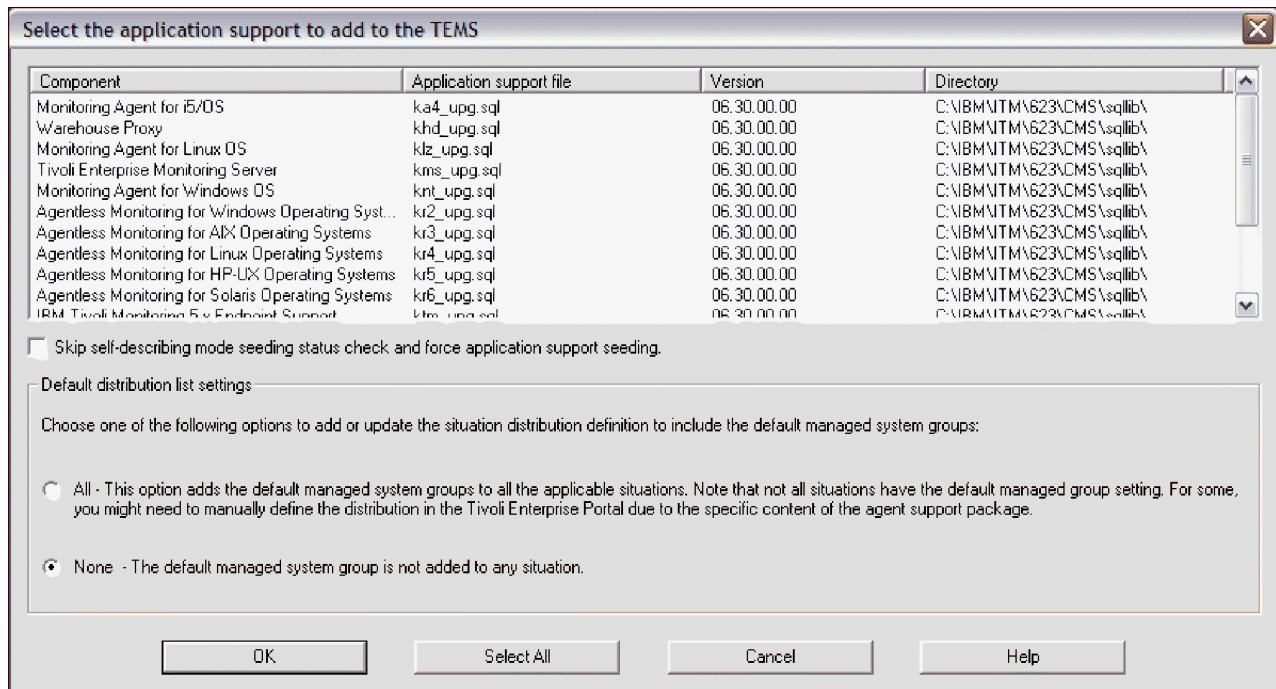


Figure 53. The Select the Application Support to Add to the TEMS window

- c. For each product-specific support package the installer checks if the Tivoli Enterprise Monitoring Server database was previously seeded with product-specific support in the self-describing mode. If so, the selected support file is excluded from the Tivoli Enterprise Monitoring Server seeding process. If you want to overwrite the support that was seeded in self-describing mode, you can select the option to **Skip self-describing mode seeding status check**.
- d. Click **OK** on the Select the Application Support to Add to the TEMS window.
This window lists the application support packages that you selected in Step 7 on page 281. Click OK to begin seeding the monitoring server (using the SQL files listed on this window). This process can take up to 20 minutes.
- e. Click **Next** on the message that provides results for the process of adding application support (see Figure 54 on page 284).

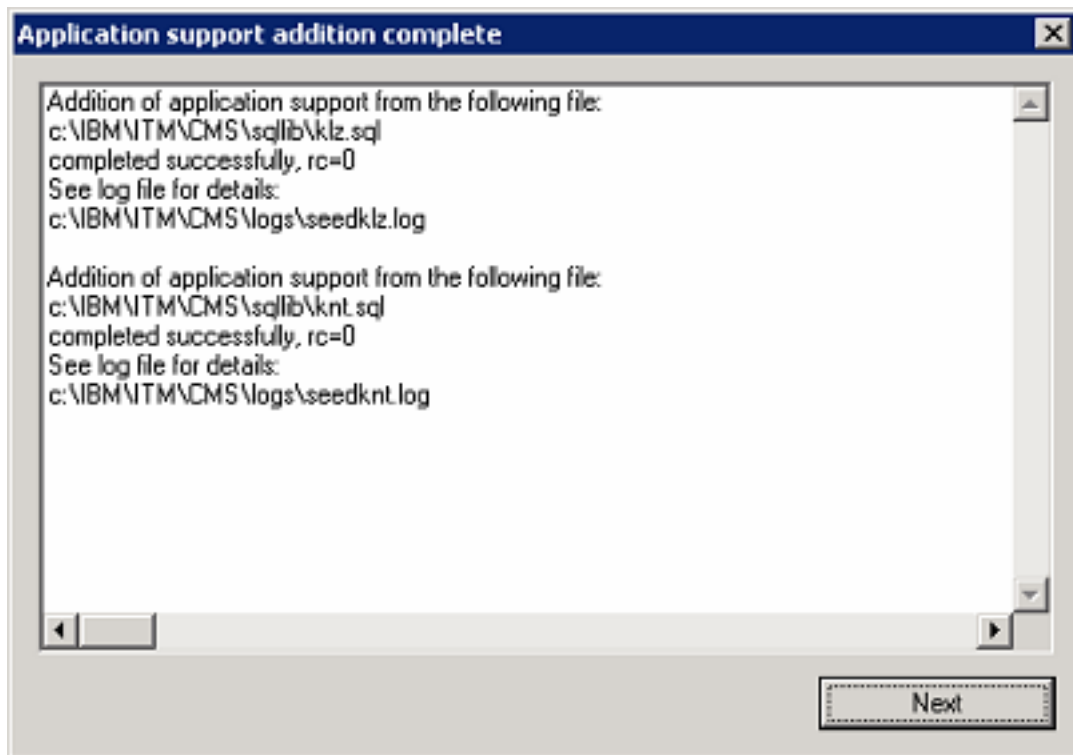


Figure 54. Application Support Addition Complete window

A return code of 0 (rc=0) indicates that the process succeeded.

Note: If the Application Support Addition Complete window is not displayed after 20 minutes, look in the IBM\ITM\CNPS\Logs\seedkpp.log files (where *pp* is the two-character code for each monitoring agent) for diagnostic messages that help you determine the cause of the problem.

13. Click **Finish** to close the installation wizard.
14. If you installed application support on the Hub monitoring server and the Tivoli Enterprise Monitoring Automation Server is also installed, restart the automation server.

Linux or UNIX: Installing application support on a monitoring server:

Complete the following steps to install application support for monitoring agents on a UNIX or Linux monitoring server.

Procedure

1. Stop the monitoring server by running the following command:
`./itmcmd server stop tems_name`

Note: If you are running in a Hot Standby environment, shut down your Hot Standby (that is, mirror) monitoring server before completing this procedure. You may restart the Hot Standby monitoring server only after you have seeded the hub server.

2. Run the following command from the installation media (the agent product CD for distributed agent products or the data files CD for z/OS agent products):
`./install.sh`
3. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory (/opt/IBM/ITM) or type the full path to the installation directory you used.
4. The following prompt is displayed:

Select one of the following:

- 1) Install products to the local host.
- 2) Exit install.

Please enter a valid number:

Enter 1 to start the installation.

The software license agreement is displayed.

5. Read through the agreement, then type 1 and press Enter to accept it.

The installer presents a list of installable components for the operating system you are currently running.

6. (*Distributed agent products only, optional*) If you are using a distributed agent product CD, optionally install monitoring agents to this computer. For example, you can install an OS to monitor the operating system. Complete the following steps if you want to install monitoring agents. If you do not want to install monitoring agents, skip to Step 8.

- a. Enter 1 to install the IBM Tivoli Monitoring components for your current operating system.
- b. Enter 1 to confirm your selection.

The installer presents a numbered list of monitoring agents that you can install.

- c. Enter the numbers of the monitoring agents that you want to install or enter the number that corresponds to All of the above. Enter more than one number on the same line separated by spaces or commas (,).

A list of the monitoring agents to be installed is displayed.

- d. Enter 1 to confirm the installation.
 - e. After the monitoring agents are installed, the installer asks if you want to install additional products or support packages. Enter 1 and go to Step 8.
7. (*z/OS agent products only, required*) If you are using a data files CD, complete the following steps to install required metaprobes. (helper programs) for the monitoring agents that you want to support.
 - a. Enter the number of the operating system. The default value is the current operating system.
 - b. Enter y to confirm your selection.

The installer presents a numbered list of monitoring agents.

- c. Enter the numbers of the monitoring agents for which you want to install application support, or enter the number that corresponds to All of the above. Enter more than one number on the same line separated by spaces or commas (,).

The installer displays the list of monitoring agents that you selected. For example:

The following products will be installed:

OMEGAMON XE for CICS on z/OS v 4.2.0
OMEGAMON XE for DB2 PE and PM on z/OS v 4.2.0 and V 5.1.0
OMEGAMON XE on z/OS v 4.2.0

Note: The prompt (The following products will be installed) seems to indicate that the installer is about to install the listed monitoring agents, which is true for distributed agents. However, for z/OS agents, only metaprobes for the monitoring agents are installed. You cannot install z/OS agents on a monitoring server on Linux or UNIX (and you cannot install monitoring agents from a data files CD).

- d. Enter 1 to confirm the installation.
 - e. After the metaprobes are installed, the installer asks if you want to install additional products or support packages. Enter y.
8. Install the application support package for the Tivoli Enterprise Monitoring Server:
 - a. Enter the number for Tivoli Enterprise Monitoring Server support.

A list of the monitoring agents for which you can install application support is displayed.

- b. Enter the numbers of the monitoring agents for which you want to install application support, or enter the number that corresponds to All of the above. Enter the numbers on the same line separated by spaces or commas (,).
- c. Enter 1 to confirm the installation.

The installation begins.

9. You are asked if you want to install application support on the Tivoli Enterprise Monitoring Server. If you reply yes, application support is automatically added.
If you disagree, you can manually add application support later; see “Installing and enabling application support” on page 275.
10. If you are updating a hub Tivoli Enterprise Monitoring Server, you are asked to choose whether you want to add the default managed system groups when you process the application-support files, as shown in Figure 33 on page 211:

All Add the default managed system groups to all applicable situations.

New Add the default managed system groups to all applicable situations from the product support packages being seeded for the first time. Modifications are not made to managed system groups in previously upgraded product support packages.

None The default managed system group is not added to any situation.

Note: Not all situations support the default managed group setting. For some, you might need to manually define the distribution using the Tivoli Enterprise Portal due to the specific content of the agent support package.

11. After you have added application support for one or more agents, you must refresh the configuration:
 - a. Start Manage Tivoli Enterprise Monitoring Services.
 - b. Pull down the Actions menu, and select the Refresh Configuration option (see Figure 55).

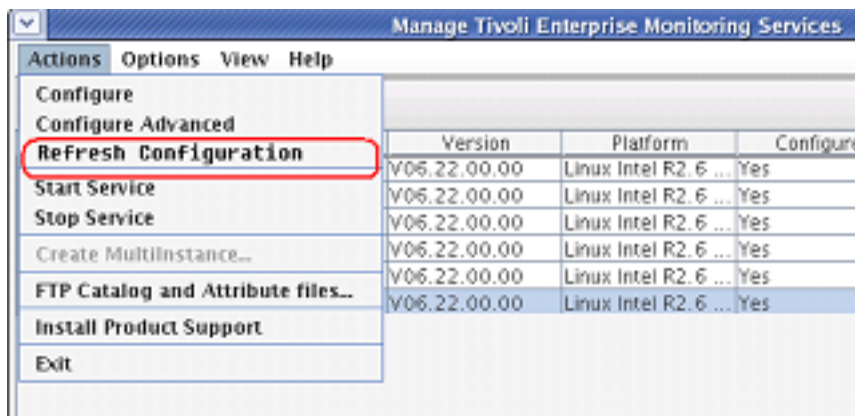


Figure 55. Refresh Configuration menu option

12. If you installed application support on the Hub monitoring server and the Tivoli Enterprise Monitoring Automation Server is also installed, restart the automation server.

Installing application support on the Tivoli Enterprise Portal Server

Use the following procedures to install application support for nonbase monitoring agents on your portal server.

- “Windows: Installing application support on a portal server” on page 287
- “Linux or AIX: Installing application support on a portal server” on page 288

Windows: Installing application support on a portal server:

Complete the following steps to install application support for monitoring agents on a Windows portal server.

Procedure

1. Stop the portal server:
 - a. Open Manage Tivoli Enterprise Monitoring Services.
 - b. Right-click **Tivoli Enterprise Portal Server** and click **Stop**.
2. In the /WINDOWS subdirectory on the agent product CD (for distributed products) or data files CD (for z/OS products), double-click the setup.exe file to launch the installation.
3. Click **Next** on the Welcome window.

Note: If a monitoring agent is already installed on this computer, select **Modify** on the Welcome window to indicate that you are updating an existing installation. Click **OK** on the message telling you about preselected items. Then skip to Step 7.

4. On the Install Prerequisites window, read the information about the required levels of IBM Global Security Toolkit (GSKit) and IBM Java.

The check box for each prerequisite is cleared if the correct level of the software is already installed. Otherwise, the check box is selected to indicate that the software is to be installed. If you are installing support from the data files CD for z/OS agent products, you might be prompted to install Oracle Java Runtime Environment (JRE) 1.4.2, even if you have already installed IBM JRE 1.5 with the distributed components of Tivoli Management Services. The two versions can coexist, and installation of application support for some monitoring agents requires Oracle Java 1.4.2. You might also see a message indicating that you can decline the installation of JRE 1.4.2 and that accepting installation of JRE 1.4.2 results in uninstallation of other Java versions. Ignore this message, because you cannot proceed without accepting the installation of Oracle Java 1.4.2, and accepting the installation does not uninstall IBM Java 1.5.

5. Click **Next**. The prerequisite software is installed if necessary.

If the installation program installs IBM GSKit or IBM JRE, you might be prompted to restart the computer when the installation is complete. If so, you receive an abort message with a **Severe error** heading. This is normal and does not indicate a problem.

If you are prompted to reboot, do the following:

- a. Click **OK** on the window prompting you to reboot.
 - b. Click **No** on the window asking whether you want to view the abort log.
 - c. Restart the computer.
 - d. Restart the installation program.
6. Click **Accept** to accept the license agreement.
 7. If you see a message regarding installed versions being newer than the agent installation, click **OK** to ignore this message.
 8. Select the application support packages that you want to install:
 - a. On the Select Features window, select **Tivoli Enterprise Portal Server**.
 - b. Expand the **Tivoli Enterprise Portal Server** node to display a list of application support packages that you can install on the portal server.

Initially, all application support packages are selected.
 - c. Clear the check boxes for application support packages that you do not want to install.

Note: If you are updating an existing installation (you selected **Modify** on the Welcome window), all check boxes on the Select Features window reflect your choices during the initial installation. Clearing a check box has the effect of *uninstalling* the component (for example, the

Tivoli Enterprise Portal Server) or product package. Clear a check box for an application support package only if you want to remove the application support.

- d. If you have other components installed on the same computer, such as the desktop client, also select those components to install the component-specific application support.

- e. Click **Next**.

- 9. On the Start Copying Files window, read the list of actions to be performed and click **Next** to start the installation.

The application support packages that you selected are installed.

- 10. On the Setup Type window, clear any components that you have already installed and configured on this computer. Click **Next**.
- 11. Type the host name for the portal server and click **Next**.
- 12. Click **Finish** to complete the installation wizard.
- 13. Restart the portal server.

Linux or AIX: Installing application support on a portal server:

Complete the following steps to install application support for monitoring agents on a Linux or AIX portal server.

Procedure

1. Stop the portal server by running the following command:
`./itmcmd agent stop cq`
2. Run the following command from the installation media (the agent product CD for distributed agent products or the data files CD for z/OS agent products):
`./install.sh`
3. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory (/opt/IBM/ITM) or enter the full path to the installation directory you used.
4. The following prompt is displayed:
Select one of the following:

1) Install products to the local host.
2) Exit install.

Please enter a valid number:

Enter 1 to start the installation.
The software license agreement is displayed.
5. Read through the agreement, then type 1 and press Enter to accept it.
The installer presents a list of installable products for the operating systems you are currently running.
6. (*Distributed agent products only, optional*) If you are using a distributed agent product CD, optionally install monitoring agents to this computer. For example, you can install an OS monitoring agent to monitor the operating system. Complete the following steps if you want to install monitoring agents. If you do not want to install monitoring agents, skip to Step 7 on page 289.
 - a. Enter 1 for your current operating system.
 - b. Enter 1 to confirm your selection.
The installer presents a numbered list of monitoring agents that you can install.
 - c. Enter the numbers of the monitoring agents that you want to install or enter the number that corresponds to All of the above. Enter more than one number on the same line separated by spaces or commas (,).
A list of the monitoring agents to be installed is displayed.

- d. Enter 1 to confirm the installation.
- e. After the monitoring agents are installed, the installer asks if you want to install additional products or support packages. Enter 1.
7. Install the application support packages for the portal server and browser client. Install application support packages for the portal desktop client if a desktop client is installed on this computer. The numbered list of items presented by the installer includes the following application support packages. (The numbers might vary from this example.)
 - 28) Tivoli Enterprise Portal Browser Client support
 - 29) Tivoli Enterprise Portal Desktop Client support
 - 30) Tivoli Enterprise Portal Server support

Note: The Tivoli Enterprise Portal Browser Client support package is portal server code that supports the browser clients. You *must* install the browser client support package on the computer where you install the portal server.

Repeat the following steps for each support package:

- a. Enter the number that corresponds to the support package (for example, 28).
A numbered list of monitoring agents is displayed.
- b. Enter the numbers that correspond to the monitoring agents for which you want to install the application support package, or enter the number that corresponds to All of the above. Type the numbers on the same line separated by spaces or commas (,).
- c. Enter 1 to confirm the installation.
The installation begins.
- d. After the support package is installed, you are asked whether you want to install additional products or product support packages. Enter 1 to install an additional package and repeat the preceding steps. Enter 2 if you are finished installing support packages.
8. Stop the portal server by running the following command:
`./itmcmd agent stop cq`
9. Run the following command to configure the portal server with the new agent information:
`./itmcmd config -A cq`

Complete the configuration as prompted. For information about configuring the portal server, see "Configuring the portal server on Linux or AIX: command-line procedure" on page 243.
10. Restart the portal server by running the following command:
`./itmcmd agent start cq`

Installing application support on the Tivoli Enterprise Portal desktop client

Use the following procedures to install application support for nonbase monitoring agents on each computer where you are running a desktop client.

Note: You must install application support on desktop clients that were installed from the installation media. You do not need to install application support on desktop clients that were obtained by using IBM Web Start for Java to download the client from the Tivoli Enterprise Portal Server.

- "Windows: Installing application support on a desktop client"
- "Linux: Installing application support on a desktop client" on page 290

Windows: Installing application support on a desktop client:

Complete the following steps to install application support for monitoring agents on a Windows desktop client.

Procedure

1. Stop the portal desktop client:
 - a. Open Manage Tivoli Enterprise Monitoring Services.
 - b. Right-click **Tivoli Enterprise Portal Desktop Client** and click **Stop**.
2. In the /WINDOWS subdirectory on the agent product CD (for distributed products) or data files CD (for z/OS products), double-click the setup.exe file to launch the installation.
3. Click **Next** on the Welcome window.

Note: If a monitoring agent is already installed on this computer, select **Modify** on the Welcome window to indicate that you are updating an existing installation. Click **OK** on the message telling you about preselected items. Then skip to Step 6.

4. Click **Next** to continue. The prerequisite software is installed if necessary.

If the installation program installs IBM GSKit or IBM JRE, you might be prompted to restart the computer when the installation is complete. If so, you receive an abort message with a **Severe error** heading. This is normal and does not indicate a problem.

If you are prompted to reboot, do the following:

- a. Click **OK** on the window prompting you to reboot.
 - b. Click **No** on the window asking whether you want to view the abort log.
 - c. Restart the computer.
 - d. Restart the installation program.
5. Read the software license agreement and click **Accept**.
 6. If you see a message regarding installed versions being newer than the agent installation, click **OK** to ignore this message.
 7. Select the application support packages that you want to install:
 - a. On the Select Features window, select **Tivoli Enterprise Portal Desktop Client**.
 - b. Expand the **Tivoli Enterprise Portal Desktop Client** node to display a list of application support packages that you can install on the portal server.
Initially, all application support packages are selected.
 - c. Clear the check boxes for application support packages that you do not want to install.

Note: If you are updating an existing installation (you selected **Modify** on the Welcome window), all check boxes on the Select Features window reflect your choices during the initial installation. Clearing a check box has the effect of *uninstalling* the component or product package. Clear a check box for an application support package only if you want to remove the application support.

- d. Click **Next**.
8. On the Start Copying Files window, read the list of actions to be performed and click **Next** to start the installation.
The application support packages that you selected are installed.
 9. On the Setup Type window, clear any components that you have already installed and configured on this computer. Click **Next**.
 10. Type the host name for the portal server and click **Next**.
 11. Click **Finish** to complete the installation wizard.

Linux: Installing application support on a desktop client:

Complete the following steps to install application support for monitoring agents on a Linux desktop client.

Before you begin

Stop the desktop client before performing this procedure.

Procedure

1. Stop the desktop client by running the following command:
`./itmcmd agent stop cj`
2. Run the following command from the installation media (the agent product CD for distributed agent products or the data files CD for z/OS agent products):
`./install.sh`
3. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory (/opt/IBM/ITM) or enter the full path to the installation directory you used.
4. The following prompt is displayed:
Select one of the following:

1) Install products to the local host.
2) Exit install.

Please enter a valid number:

Enter 1 to start the installation.
5. Read the software license agreement; then type 1 and press Enter to accept it.
The installer presents a list of installable components for your current operating system.
6. (*Distributed agent products only, optional*) If you are using a distributed agent product CD, optionally install monitoring agents to this computer. For example, you can install an OS monitoring agent to monitor the operating system. Complete the following steps if you want to install monitoring agents. If you do not want to install monitoring agents, skip to Step 7.
 - a. Enter 1 for your current operating system.
 - b. Enter 1 to confirm your selection.
The installer presents a numbered list of monitoring agents that you can install.
 - c. Enter the numbers of the monitoring agents that you want to install or enter the number that corresponds to All of the above. Enter more than one number on the same line separated by spaces or commas (,).
A list of the monitoring agents to be installed is displayed.
 - d. Enter 1 to confirm the installation.
 - e. After the monitoring agents are installed, the installer asks if you want to install additional products or support packages. Enter 1.
7. Install the application support package for the portal desktop client:
 - a. Enter the number that corresponds to Tivoli Enterprise Portal Desktop Client support.
A numbered list of monitoring agents is displayed.
 - b. Enter the numbers that correspond to the monitoring agents for which you want to install the application support package, or enter the number that corresponds to All of the above. Type the numbers on the same line separated by spaces or commas (,).
 - c. Enter 1 to confirm the installation.
The installation begins.
8. After application support for all monitoring agents is installed, you are asked whether you want to install additional products or product support packages. Enter 2.
9. Run the following command to configure the desktop client with the new agent information:
`./itmcmd config -A cj`

Complete the configuration as prompted. For information about configuring the desktop client, see “Linux: Configuring the desktop client” on page 275.

- Restart the desktop client by running the following command:

```
./itmcmd agent start cj
```

Configuring application support on nonlocal monitoring servers

The following sections contain procedures for installing and enabling application support on a monitoring server (hub or remote) that is located on a remote computer. For example, you might need to install application support on a hub monitoring server on z/OS for a monitoring agent running on Windows or a Linux operating system. Or you might want to install application support on a monitoring server or client on a different computer than the one on which you are installing an agent or a portal server.

If you are using a z/OS hub and you want to aggregate and prune historical data for any agent, you must transfer the catalog and attribute files for the Summarization and Pruning Agent to the hub.

The procedures in this section require that either a Tivoli Enterprise Portal Server or a Tivoli Enterprise Monitoring Server be installed on the Windows, Linux, or UNIX computer and configured to communicate with the monitoring server on the nonlocal computer.

- “Configuring application support on a nonlocal monitoring server from a Windows system”
- “Configuring application support on a nonlocal monitoring server from a Linux or UNIX system” on page 294

Configuring application support on a nonlocal monitoring server from a Windows system

This section describes how you can install and enable application support on a nonlocal monitoring server from your local Windows system. The nonlocal monitoring server can be a hub or remote monitoring server installed on a Windows, Linux, UNIX, or z/OS computer. A monitoring server or a portal server must be installed on the local computer and application support for the agent or agents must be installed on it.

You add support to a nonlocal monitoring server by copying the *kpp.cat* and *kpp.atr* files for each from the local computer to the nonlocal computer. If you are adding support to a nonlocal hub monitoring server, you use Manage Tivoli Enterprise Monitoring Services to transfer the *kpp.sql* files for each agent.

Copying the CAT and ATR files to the nonlocal monitoring server:

Copy the .cat and .atr files for the agents you want to support from the local Windows monitoring server to the nonlocal monitoring server. If you use ftp, copy the files in ASCII format.

The .cat and .atr files are located in the following directories on the local monitoring server:

- .cat files are located in *install_dir*\cms\RKDSCATL
- .atr files are located in *install_dir*\cms\ATTRLIB

Copy the files to the following directories on the remote computer:

Table 55. Locations of CAT and ATR files for the monitoring server

Remote computer on:	File type	Directory
Windows	.cat	<i>install_dir</i> \cms\RKDSCATL
	.atr	<i>install_dir</i> \cms\ATTRLIB
Linux UNIX	.cat	<i>install_dir</i> /tables/cicatrsql/RKDSCATL
	.atr	<i>install_dir</i> /tables/cicatrsql/ATTRLIB

where *install_dir* specifies the IBM Tivoli Monitoring installation directory. The IBM Tivoli Monitoring installation directory is represented by the %CANDLE_HOME% (Windows) or \$CANDLEHOME (Linux and UNIX) environment variable. The default installation directory on Windows is \IBM\ITM. The default installation directory on Linux and UNIX is /opt/IBM/ITM.

Note:

1. If you **export** the CANDLEHOME environment variable to your current session, many of the installation and configuration commands do not require that CANDLEHOME be passed to them (usually via the **-h** CLI parameter).
2. If you are adding support to a monitoring server on z/OS, you can use the FTP utility provided with Manage Tivoli Enterprise Monitoring Services. While checking existing files on the z/OS machine, the tool extracts the version line from the content of all files found. In a CAT file, the version line begins with the character @. In an ATR file, the version line begins with the comment marks //. The version numbers and timestamps that are obtained are used for version comparison. The following rules apply:
 - Files selected for FTP transfer that have a version higher than the corresponding file on z/OS, or have no corresponding file on z/OS, are transferred.
 - Files selected for FTP transfer that have a version lower than the corresponding file on z/OS are not transferred. You are alerted that some files cannot be transferred due to possible backleveling, and those files are excluded from the list of files to be transferred.
 - Files selected for FTP transfer for which it was impossible to determine a version are not transferred. This can be caused by, for example, file access error, or version text not found or incorrect. You are alerted that some files cannot be transferred due to the encountered errors, and those files are excluded from the list of files to be transferred.
 - If all files selected for FTP transfer are excluded from the transfer due to possible backleveling or reading version errors, the transfer is cancelled.
 - If some files selected for FTP transfer are excluded from the transfer due to possible backleveling or reading version errors, you are alerted that some files cannot be transferred and those files are excluded from the list of files to be transferred. The remaining files are transferred after confirmation.

For more information, see *Configuring the Tivoli Enterprise Monitoring Server on z/OS* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm).

3. If you specify an incorrect directory name, you will receive the following error:

The IBM Tivoli Monitoring installation directory cannot exceed 80 characters or contain non-ASCII, special or double-byte characters.

The directory name can contain only these characters:

"abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ _:0123456789()~./".

Adding application support (SQL files) to a nonlocal hub:

If you are adding application support to a nonlocal hub monitoring server, you must add the SQL files used by the EIB, in addition to the catalog and attribute files.

About this task

Use the following procedure to seed the hub with SQL files using the Manage Tivoli Enterprise Monitoring Services utility on your local Windows system:

Procedure

1. Ensure that the hub monitoring server is started.

Note: If you are running in a Hot Standby environment, shut down your Hot Standby (that is, mirror) monitoring server before completing this procedure. You may restart the Hot Standby monitoring server only after you have seeded the hub server.

2. In the Manage Tivoli Enterprise Monitoring Services window, select **Actions → Advanced → Add TEMS application support**.
3. On the Add Application Support to the TEMS window, Select **On a different computer** and click **OK**.
4. When you are prompted to ensure that the hub monitoring server is configured and running, click **OK**.
5. On the Non-Resident TEMS Connection window, provide the hub monitoring server TEMS name (node ID) and select the communication protocol to use in sending the application support SQL files to the hub monitoring server.
6. On the next window, provide any values required by the communication protocol.
For example, if the protocol is IP.PIPE, you are prompted for the fully qualified TCP/IP host name and port number of the z/OS system where the monitoring server is installed. See for the values you recorded during installation planning.
7. On the Select the Application Support to Add to the TEMS window, select the products for which you want to add application support or click Select All to choose all available products. Click **OK**.
The SQL application support files are added to the hub monitoring server. This might take several minutes.
8. The Application Support Addition Complete window shown in Figure 54 on page 284 gives you information about the status and location of the application support SQL files. Click **Save As** if you want to save the information in a text file. Click Close to close the window.
If the Application Support Addition Complete window is not displayed after 20 minutes, look in the IBM\ITM\CNPS\Logs\seedkpp.log files (where pp is the two-character code for each monitoring agent) for diagnostic messages that help you determine the cause of the problem.
9. If the monitoring server is not already stopped, stop it.
10. Restart the monitoring server.

Configuring application support on a nonlocal monitoring server from a Linux or UNIX system

This section describes how you can install and enable application support on a nonlocal monitoring server from your local Linux or UNIX system. The nonlocal monitoring server can be a hub or remote monitoring server installed on a Windows, Linux, or UNIX system.

Before you begin:

- As a prerequisite, you must install a monitoring server, a portal server or a monitoring agent on the local Linux or UNIX computer. This step is necessary to make the Manage Tivoli Enterprise Monitoring Services available on the local computer. If you do not have a monitoring server, you must install the support files on the local computer using the procedure described in “Installing application support files on a computer with no monitoring server” on page 298.
- Determine which monitoring agents you want to support on the nonlocal monitoring server. Ensure that the application support files (.cat, .atr, and .sql files) for those agents are available on the local monitoring server:
 - Application support files are located in the following directories on a Linux or UNIX monitoring server:

Table 56. Locations of application support files on a Linux or UNIX monitoring server

File type	Directory
.cat	<i>install_dir</i> /tables/cicatrsq/RKDSCATL
.atr	<i>install_dir</i> /tables/cicatrsq/ATTRLIB

Table 56. Locations of application support files on a Linux or UNIX monitoring server (continued)

File type	Directory
.sql	<i>install_dir</i> /tables/cicatsrq/SQLLIB

- The file names of the application support files have the following format:
`kpc.ext`
 where *pc* is the product code for the agent and *ext* is the file extension.
 For example, kud.sql is the SQL support file for the DB2 for Linux, UNIX, and Windows monitoring agent. See Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883 for a list of product codes.
- If you cannot find application support files for some agents for which you want to install application support, install the missing files on this computer.
 - To install missing support files for base monitoring agents, follow the installation steps described in “Installing the monitoring server” on page 201.
 - To install missing files for nonbase monitoring agents, follow the installation steps described in “Linux or UNIX: Installing application support on a monitoring server” on page 284.
 - If no monitoring server is installed on this computer, use the procedure in “Installing application support files on a computer with no monitoring server” on page 298.
- Gather the following information about the monitoring server on the remote computer:
 - The host name or IP address.
 - The protocol and port number that was specified when the monitoring server was configured.
 The monitoring server on the remote computer must be configured to use the IP.UDP, IP.PIPE, or IP.SPIPE communications protocol. This procedure does not support a monitoring server that was configured to use SNA.
- Verify that the monitoring server on the remote computer is running.
- Verify that the hub monitoring server to which this remote server sends its data is running.
- In these instructions, *install_dir* specifies the IBM Tivoli Monitoring installation directory. You can enter either \$CANDLEHOME or the name of the directory. The default installation directory on Linux and UNIX is /opt/IBM/ITM.
- If you are running in a Hot Standby environment, shut down your Hot Standby (that is, mirror) monitoring server before completing this procedure. You may restart the Hot Standby monitoring server only after you have seeded the hub server.

Copying the CAT and ATR files to the nonlocal monitoring server: Copy the .cat and .atr files for the agents you want to support from the local Linux or UNIX monitoring server to the nonlocal monitoring server. If you use FTP, copy the files in ASCII format. The .cat and .atr files are located in the following directories on the local monitoring server:

- CAT files are located in *install_dir*/tables/cicatsrq/RKDSCATL
- ATR files are located in *install_dir*/tables/cicatsrq/ATTRLIB

Copy the files to the directory shown in Table 57 on the remote computer:

Table 57. Locations of CAT and ATR files for the monitoring server

Remote computer on:	File type	Directory
Windows	.cat	<i>install_dir</i> \cms\RKDSCATL
	.atr	<i>install_dir</i> \cms\ATTRLIB
Linux UNIX	.cat	<i>install_dir</i> /tables/cicatsrq/RKDSCATL
	.atr	<i>install_dir</i> /tables/cicatsrq/ATTRLIB

where *install_dir* specifies the IBM Tivoli Monitoring installation directory. The IBM Tivoli Monitoring installation directory is represented by the %CANDLE_HOME% (Windows) or \$CANDLEHOME (Linux and UNIX) environment variable. The default installation directory on Windows is \IBM\ITM. The default installation directory on Linux and UNIX is /opt/IBM/ITM.

Note:

1. If you are adding support to a monitoring server on z/OS, you can use the FTP utility provided with Manage Tivoli Enterprise Monitoring Services. While checking existing files on the z/OS machine, the tool extracts the version line from the content of all files found. In a CAT file, the version line begins with the character @. In an ATR file, the version line begins with the comment marks //. The version numbers and timestamps that are obtained are used for version comparison. The following rules apply:
 - Files selected for FTP transfer that have a version higher than the corresponding file on z/OS, or have no corresponding file on z/OS, are transferred.
 - Files selected for FTP transfer that have a version lower than the corresponding file on z/OS are not transferred. You are alerted that some files cannot be transferred due to possible backleveling, and those files are excluded from the list of files to be transferred.
 - Files selected for FTP transfer for which it was impossible to determine a version are not transferred. This can be caused by, for example, file access error, or version text not found or incorrect. You are alerted that some files cannot be transferred due to the encountered errors, and those files are excluded from the list of files to be transferred.
 - If all files selected for FTP transfer are excluded from the transfer due to possible backleveling or reading version errors, the transfer is cancelled.
 - If some files selected for FTP transfer are excluded from the transfer due to possible backleveling or reading version errors, you are alerted that some files cannot be transferred and those files are excluded from the list of files to be transferred. The remaining files are transferred after confirmation.

For more information, see *Configuring the Tivoli Enterprise Monitoring Server on z/OS* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm).

2. If you specify an incorrect directory name, you will receive the following error:

The IBM Tivoli Monitoring installation directory cannot exceed 80 characters or contain non-ASCII, special or double-byte characters.

The directory name can contain only these characters:

"abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ _\.:0123456789()~-./*".

If the nonlocal monitoring server to which you are adding application support is a hub, you must also install the SQL files used by the EIB. See "Adding application support (SQL files) to a nonlocal hub."

Adding application support (SQL files) to a nonlocal hub:

To add application support SQL files to a hub monitoring server on a nonlocal system, complete this procedure.

Procedure

1. Enable the GUI interface.

Your Linux or UNIX environment might already have a GUI interface enabled. Otherwise, perform the following tasks to enable it:

 - a. Enable X11.
 - b. Make sure you have access to a native X-term monitor or an X-Window emulator.
 - c. If using an X-Window emulator, enable X11 access to the X-Window server (command example: `xhost +`).

- d. If using an X-Window emulator, set the display environment variable to point to the X-Window server:

```
export DISPLAY=pc_ip_address:0
```
2. Ensure that the hub monitoring server is running.
3. To start Manage Tivoli Enterprise Monitoring Services, go to the \$CANDLEHOME bin directory (example: /opt/IBM/ITM/bin) and run this command:

```
./itmcmd manage &
```

A GUI window opens for Manage Tivoli Enterprise Monitoring Services.
4. Select **Actions** → **Install Product Support**.

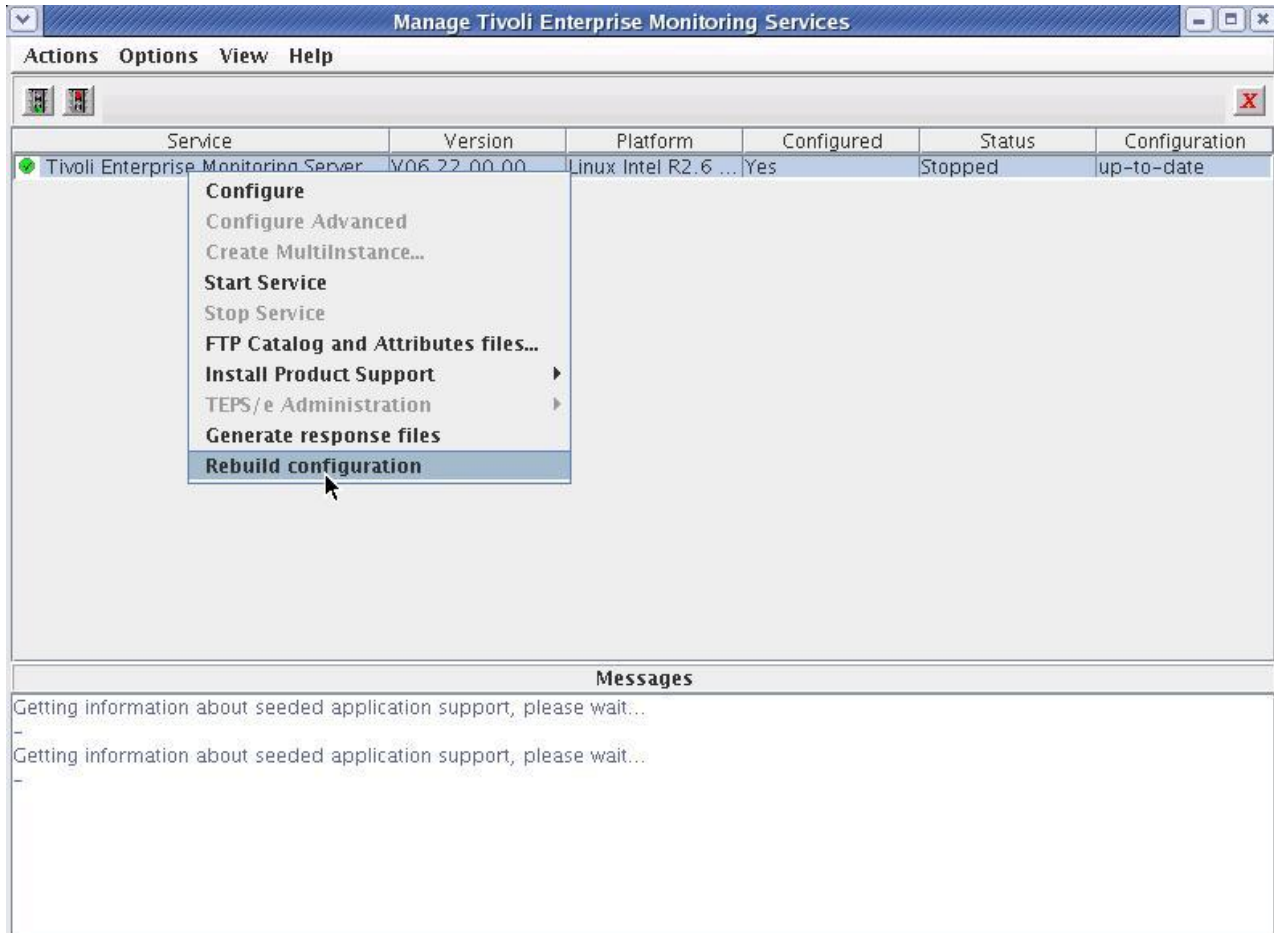


Figure 56. Manage Tivoli Enterprise Monitoring Services Install Product Support window

5. On the Add Application Support to the TEMS window, select **On a different computer** and click **OK**.
6. When you are prompted to ensure that the hub monitoring server is configured and running, click **OK**.
7. On the Non-Resident TEMS Connection window, provide the hub monitoring server name (node ID) and select the communication protocol to use in sending the application support SQL files to the hub monitoring server.
8. Select the appropriate communications protocol and click **OK**.
9. On the next window, supply any values required by the selected communication protocol and click **OK**.

10. On the Install Product Support window, select the monitoring agents for which you want to add application support to the hub monitoring server, and click Install.
11. In Manage Tivoli Enterprise Monitoring Services, look for this message:
Remote seeding complete!
12. Stop and restart the hub monitoring server.

Installing application support files on a computer with no monitoring server:

You can install application support files for a monitoring server on a UNIX or Linux computer where no monitoring server is installed and then use the files to configure support on a nonlocal monitoring server.

About this task

Use the following procedure to install application support files on computer on which no monitoring server is installed for use on a monitoring server on another computer:

Procedure

1. Run the following command from the installation media, either the IBM Tivoli Monitoring base media, the agent product CD for distributed agent products or the data files CD for z/OS agent products:
./install.sh
2. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory (/opt/IBM/ITM) or enter the full path to the installation directory you used.
The following prompt is displayed:
Select one of the following:
 - 1) Install products to the local host.
 - 2) Install products to depot for remote deployment (requires TEMS).
 - 3) Install TEMS support for remote seeding
 - 4) Exit install.
Please enter a valid number:
3. Type 3 and press Enter.
The software license agreement is displayed.
4. Read through the software license agreement, then type 1 and press Enter to accept the agreement.
The installer presents a list of currently installed products, the operating systems for which product packages are available, and component support categories.
5. Type the 32-character encryption key to use and press Enter, or press enter to use the default key. This key must be the same key as that used during the installation of the portal server to which the client will connect.
The installer presents a list of products for which support packages are available.
6. Enter the numbers of the monitoring agents that you want to install or enter the number that corresponds to All of the above. Enter more than one number on the same line separated by spaces or commas (,).
A list of the monitoring agents to be installed is displayed.
7. Confirm the selection by pressing Enter.
The installer asks if you want to install additional products or support packages.
8. Type 2 and press Enter.
The support packages are installed. When installation is complete, you will see the message
... finished postprocessing.
Installation step complete.

Using IBM Installation Manager to install or update IBM Tivoli Monitoring components for a dashboard environment

The IBM Installation Manager simplifies the procedure for installing or updating IBM Tivoli Monitoring components for a dashboard environment.

Install the Jazz for Service Management components that are described in “Required software and memory requirements for a dashboard environment” on page 91 before you install the IBM Infrastructure Management Dashboards for Servers, Tivoli Authorization Policy Server, and tivcmd Command Line Interface for Authorization Policy components,

Note: If you use an IBM Installation Manager network repository or local repositories, you can install the Jazz for Service Management components and the IBM Tivoli Monitoring dashboard components at the same time. IBM Installation Manager ensures that the components are installed in the correct order. See the Jazz for Service Management Installation Guide for more details on using IBM Installation Manager repositories to install Jazz for Service Management, go to the Jazz for Service Management Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/psc_ic-homepage.html). The procedures below outline how to create or update a network or local based repository for the IBM Tivoli Monitoring components.

Choose one of the following methods to install or update the IBM Infrastructure Management Dashboards for Servers, Tivoli Authorization Policy Server, and tivcmd Command Line Interface for Authorization Policy components:

- “Using the Installation Launchpad to install or update IBM Tivoli Monitoring components for a dashboard environment”
- “Using IBM Installation Manager to install or update components from a local repository” on page 300
- “Using IBM Installation Manager to install or update components from a network repository” on page 302
- “Saving package information for rollback” on page 305

Note: The Installation Launchpad is the recommended method if you do not have experience of using IBM Installation Manager. For more background on using local or network repositories, see “Using IBM Installation Manager” on page 42.

You can also configure IBM Installation Manager to save package information so that you can perform a rollback after installing an updated package in the future. For more information, see “Saving package information for rollback” on page 305.

Using the Installation Launchpad to install or update IBM Tivoli Monitoring components for a dashboard environment

The IBM Tivoli Monitoring Installation Launchpad simplifies the procedure for installing or updating the IBM Infrastructure Management Dashboards for Servers, Tivoli Authorization Policy Server, and Tivoli Policy Command-Line Interface. The Installation Launchpad automatically installs IBM Installation Manager version 1.6.1 if it is not installed on the system, or if it is at an earlier version. The Installation Launchpad also configures an IBM Installation Manager local repository with the location of the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image, and uses the IBM Installation Manager graphical user interface to guide you through the installation or update of one or more of the IBM Tivoli Monitoring components for a dashboard environment.

When you use the IBM Tivoli Monitoring Installation Launchpad to install or update the IBM Infrastructure Management Dashboards for Servers and Tivoli Authorization Policy Server, IBM Dashboard Application Services Hub must already be installed on the system.

For a complete list of prerequisite criteria, see “Required software and memory requirements for a dashboard environment” on page 91.

For more information about using the IBM Tivoli Monitoring Installation Launchpad to install or update components for a dashboard environment, see the following topics:

- “Installing the IBM Infrastructure Management Dashboards for Servers using the graphical user interface” on page 308
- “Updating the IBM Infrastructure Management Dashboards for Servers using the graphical user interface” on page 315
- “Installing and configuring the Tivoli Authorization Policy Server using the graphical user interface” on page 320
- “Installing the tivcmd Command Line Interface for Authorization Policy using the graphical user interface” on page 330

Using IBM Installation Manager to install or update components from a local repository

With IBM Installation Manager you can create a local repository that specifies where to find the components to install or update. The repository can be on the local file system or on connected installation media.

You then use one of the following IBM Installation Manager modes to install or update the components in the repository:

- Wizard mode, which uses a graphical user interface.
- Console mode, which allows you to perform the installation or update from a command line.
- Silent installation mode, which uses a response file to determine what to install or update.

For more information about IBM Installation Manager and supported repository types, see <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.

You can install or update one or more of the IBM Tivoli Monitoring components for a dashboard environment from the same local repository. The topics that follow this section outline how to use the graphical user interface and console mode to create a local repository for the components. For details on the silent installation mode, see “Performing a silent installation of software packages using IBM Installation Manager” on page 855.

You can also set up a separate local repository for the Jazz for Service Management components if you want to use IBM Installation Manager to install or update those components instead of using the Jazz for Service Management installation launchpad application. Then, you can use IBM Installation Manager to install or update the Jazz for Service Management and IBM Tivoli Monitoring components for a dashboard environment at the same time.

Note:

1. If you use IBM Installation Manager to install Jazz for Service Management, you cannot install the Tivoli Common Reporting component of Jazz for Service Management. You must use the Jazz for Service Management installation launchpad to install the Tivoli Common Reporting component. For more details about using IBM Installation Manager repositories to install Jazz for Service Management, see the Jazz for Service Management Installation Guide at http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html.
2. You cannot use IBM Installation Manager console mode to install Jazz for Service Management. However, the graphical user interface and silent mode are supported.

Creating a local repository connection using the IBM Installation Manager graphical user interface

Follow this procedure to use the IBM Installation Manager graphical user interface to create a local repository connection that contains the IBM Tivoli Monitoring components for a dashboard environment.

Before you begin

Obtain the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image.

About this task

Follow this procedure to install the components from a local repository:

Procedure

1. Install IBM Installation Manager from the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image, if it is not already installed or is at a version before V1.6.1. IBM Installation Manager version 1.6.1 is provided on the DVD under the *im.<operating system>.<architecture>* directories. For more details about the installation, see the IBM Installation Manager information center at <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.
2. Start the IBM Installation Manager in graphical user interface mode:
 - **Windows** Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager.
 - **Linux** **UNIX** Execute the IBMIM binary under /opt/IBM/InstallationManager/eclipse.
3. Click **File** → **Preferences**.
4. Click **Add Repository**.
5. In the **Repository** field, browse to the location of the diskTag.inf file on the *Dashboards for Servers and Authorization Policy Components* DVD or DVD image and click **OK**.
6. The diskTag.inf file location should now be listed on the **Repository Preferences** panel. Select the **Location** check box and click **OK**. This tests the connection to the local repository.

What to do next

Use IBM Installation Manager to install or update the IBM Tivoli Monitoring components for a dashboard environment, as described in the following sections:

- “Installing and configuring the IBM Infrastructure Management Dashboards for Servers” on page 306
- “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318

Creating a local repository connection using the IBM Installation Manager console mode

Follow this procedure to use the IBM Installation Manager console mode to create a local repository connection that contains the IBM Tivoli Monitoring components for a dashboard environment.

Procedure

1. Install IBM Installation Manager from the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image, if it is not already installed or is at a version before V1.6.1. IBM Installation Manager version 1.6.1 is provided on the DVD under the *im.<operating system>.<architecture>* directories. For more details about the installation, see the IBM Installation Manager information center at <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.
2. Start the IBM Installation Manager in console mode from the directory where it was installed. For example:

- **Windows** C:\Program Files\IBM\Installation Manager\eclipse\tools\imcl.exe -c
- **Linux** **UNIX** /opt/IBM/InstallationManager/eclipse/tools/imcl -c

3. The following menu is displayed:

```
Select:
  1. Install - Install software packages
  2. Update - Find and install updates and fixes to installed software packages
  3. Modify - Change installed software packages
  4. Roll Back - Revert to an earlier version of installed software packages
  5. Uninstall - Remove installed software packages

Other Options:
  L. View Logs
  S. View Installation History
  V. View Installed Packages
  -----
  P. Preferences
  -----
  E. Export Data for Problem Analysis
  A. About IBM Installation Manager
  -----
  X. Exit Installation Manager
```

Enter P: Preferences

4. Enter 1: Repositories.
5. Enter D: Add repository.
6. Enter the location of the diskTag.inf file on the *Dashboards for Servers and Authorization Policy Components* DVD or DVD image.

If you add a repository that requires credentials, you are prompted to provide the required credentials. Use the correct case when you enter the repository location. If the correct case is not used, the package is not shown in the list of available packages for installation.

- a. Enter P: Provide credentials and connect.
 - b. Enter the user_name and press Enter.
 - c. Enter the password and press Enter.
 - d. Enter 1 to save the password.
 - e. Enter 0: Ok
7. Enter A: Apply Changes and Return to Preferences menu.
 8. Enter R: Return to Main menu.

What to do next

Use IBM Installation Manager to install or update the IBM Tivoli Monitoring components for a dashboard environment, as described in the following sections:

- “Installing and configuring the IBM Infrastructure Management Dashboards for Servers” on page 306
- “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318

Using IBM Installation Manager to install or update components from a network repository

With IBM Installation Manager you can set up a central repository that contains one or more installation or update packages. This repository is created by using the IBM Packaging Utility and is accessed using HTTP from each of the computer systems where you install the IBM Tivoli Monitoring components for a dashboard environment.

For more information about the IBM Packaging Utility, go to http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp?topic=%2Fcom.ibm.cic.auth.ui.doc%2Fhelpindex_pu.html.

On each computer system where you want to install or update one of the IBM Tivoli Monitoring components for a dashboard environment, you use one of the following IBM Installation Manager modes to configure the location of the network repository and install or update the components in the repository:

- Wizard mode, which uses a graphical user interface.
- Console mode, which allows you to perform the installation or update from a command line.
- Silent installation mode, which uses a response file to determine what to install or update.

For more information about IBM Installation Manager and supported repository types, see <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.

You can install or update one or more of the IBM Tivoli Monitoring components for a dashboard environment from the same network repository. The topics that follow this section provide information about using the graphical user interface and console modes to use a network repository to install and configure these components. For details on the silent installation mode, see “Performing a silent installation of software packages using IBM Installation Manager” on page 855.

You can also set up the network repository to include the Jazz for Service Management components if you want to use IBM Installation Manager to install or update them from a central HTTP server instead of using the Jazz for Service Management installation launchpad application. Then, you can use IBM Installation Manager to install or update the Jazz for Service Management and IBM Tivoli Monitoring components for a dashboard environment at the same time.

Note:

1. If you use IBM Installation Manager to install Jazz for Service Management, you cannot install the Tivoli Common Reporting component of Jazz for Service Management. You must use the Jazz for Service Management installation launchpad to install the Tivoli Common Reporting component. For the Jazz for Service Management Installation Guide and more details on using IBM Installation Manager repositories to install Jazz for Service Management, go to http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html.
2. IBM Installation Manager console mode cannot be used to install Jazz for Service Management. However, the graphical user interface and silent mode are supported.


Creating a network repository connection using the IBM Installation Manager graphical user interface

Follow this procedure to use the IBM Installation Manager graphical user interface to create a connection to the network repository that contains the dashboard environment components. Complete this procedure on each computer system where you are installing or updating one of the IBM Tivoli Monitoring components for a dashboard environment.

Before you begin

Use the IBM Packaging Utility to create a network repository that contains the IBM Tivoli Monitoring components on the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD.

Procedure

1. Install IBM Installation Manager from the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image, if it is not already installed or is at a version before v1.6.1. IBM Installation Manager version 1.6.1 is provided on the DVD under the *im.<operating system>.<architecture>* directories. For more details about the installation, see the IBM Installation Manager information center at <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.
2. Start the IBM Installation Manager in graphical user interface mode:
 -  Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager.

- **Linux** **UNIX** Execute the IBMIM binary under /opt/IBM/InstallationManager/eclipse.
3. Click **File** → **Preferences**.
 4. Click **Add Repository**.
 5. In the **Repository** field, enter the HTTP URL of the network repository and click **OK**.
 6. The network repository location should now be listed on the **Repository Preferences** panel. Select the **Location** check box and click **OK**. This tests the connection to the network repository.

What to do next

Use IBM Installation Manager to install or update the IBM Tivoli Monitoring components for a dashboard environment, as described in the following sections:

- “Installing and configuring the IBM Infrastructure Management Dashboards for Servers” on page 306
- “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318

Creating a network repository connection using the IBM Installation Manager console mode

Follow this procedure to use the IBM Installation Manager console mode to create a network repository connection that contains the IBM Tivoli Monitoring components for a dashboard environment.

Before you begin

Use the IBM Packaging Utility to create a network repository that contains the IBM Tivoli Monitoring components on the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components DVD*.

Procedure

1. Install IBM Installation Manager from the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components DVD* or DVD image, if it is not already installed or is at a version before V1.6.1. IBM Installation Manager version 1.6.1 is provided on the DVD under the *im.<operating system>.<architecture>* directories. For more details about the installation, see the IBM Installation Manager information center at <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.
2. Start the IBM Installation Manager in console mode from the directory where it was installed. For example:
 - **Windows** C:\Program Files\IBM\Installation Manager\eclipse\tools\imcl.exe -c
 - **Linux** **UNIX** /opt/IBM/InstallationManager/eclipse/tools/imcl -c
3. The following menu is displayed:

```
Select:
  1. Install - Install software packages
  2. Update - Find and install updates and fixes to installed software packages
  3. Modify - Change installed software packages
  4. Roll Back - Revert to an earlier version of installed software packages
  5. Uninstall - Remove installed software packages

Other Options:
  L. View Logs
  S. View Installation History
  V. View Installed Packages
  -----
  P. Preferences
  -----
  E. Export Data for Problem Analysis
  A. About IBM Installation Manager
  -----
  X. Exit Installation Manager
```

Enter P: Preferences

4. Enter 1: Repositories.
5. Enter D: Add repository.
6. Enter the HTTP URL of the network repository.

If you add a repository that requires credentials, you are prompted to provide the required credentials. Use the correct case when you enter the repository location. If the correct case is not used, the package is not shown in the list of available packages for installation.

- a. Enter P: Provide credentials and connect.
 - b. Enter the user_name and press Enter.
 - c. Enter the password and press Enter.
 - d. Enter 1 to save the password.
 - e. Enter 0: Ok
7. Enter A: Apply Changes and Return to Preferences menu.
 8. Enter R: Return to Main menu.

What to do next

Use IBM Installation Manager to install or update the IBM Tivoli Monitoring components for a dashboard environment, as described in the following sections:

- “Installing and configuring the IBM Infrastructure Management Dashboards for Servers” on page 306
- “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318




Saving package information for rollback

You can use the IBM Installation Manager rollback feature to revert an updated package back to a previous version for those packages that support rollback.

About this task

To complete the rollback, the Installation Manager requires information from the earlier version of the package for those packages that support rollback. By default, the files that are needed to complete the rollback are saved on your computer when each version of the package is installed. On the **Files for Rollback** preference page, you can change the default action and remove any previously stored files.

Procedure

1. Start the IBM Installation Manager:
 -  Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager\eclipse.
 -   Execute the IBMIM binary under /opt/IBM/InstallationManager/eclipse.
2. Click **File → Preferences → Files for Rollback**.
3. Select **Save file for rollback** to save on your computer a copy of files that are needed to rollback packages.
4. Click **Delete Saved Files** to delete all saved files. If you delete the files and you need to rollback a package, you must connect to your old repository or insert the media for the previous version of the package to obtain the required files.

Installing and configuring the IBM Infrastructure Management Dashboards for Servers

IBM Infrastructure Management Dashboards for Servers allows IBM Dashboard Application Services Hub users to display events for monitoring agents, managed system groups, and key Linux OS agent, UNIX OS agent, and Windows OS agent health metrics.

For information about the software that you must install before you can install the infrastructure management dashboards application, see “Required software and memory requirements for a dashboard environment” on page 91.

Note: IBM Infrastructure Management Dashboards for Servers supports only managed system groups and managed systems from a single Hub monitoring server. If you have multiple Hub monitoring servers, you must install an instance of IBM Dashboard Application Services Hub and IBM Infrastructure Management Dashboards for Servers for each IBM Tivoli Monitoring environment. For more information, see “Multi-hub environments” on page 63.

The IBM Infrastructure Management Dashboards for Servers package has two features: **Installation** and **Configuration** that you can select during installation and updates. When you choose the Installation feature, the infrastructure management dashboards files are copied from the installation media or repository to the file system. When you choose the Configuration feature, the infrastructure management dashboards application is deployed into IBM Dashboard Application Services Hub using the scripts that were installed by the Installation feature. Normally you should select both features during installation. However, if you want to view the deployment scripts before they are executed, you can just choose the Installation feature during installation. After reviewing the deployment scripts, you use the Modify operation of IBM Installation Manager to select the Configuration feature and deploy the infrastructure management dashboards application into IBM Dashboard Application Services Hub.

Note: Although you could use the WebSphere Application Server wsadmin command-line interface to run the Infrastructure Management Dashboard for Servers deployment scripts, this method is not supported. IBM Installation Manager must be used to track which features are installed and configured.

See the following table for tasks that you perform to install, update, and configure Infrastructure Management Dashboards for Servers:

Table 58. Installation, update, and configuration topics

Task	Where to find detailed information
<p>Install and configure by using the IBM Tivoli Monitoring installation launchpad or the IBM Installation Manager graphical user interface.</p> <p>Note:</p> <ol style="list-style-type: none"> 1. The Tivoli Monitoring installation launchpad is the recommended approach if you are not familiar with IBM Installation Manager and you have a supported browser on the computer system where you are installing. 2. You can install one or more of the IBM Tivoli Monitoring packages for a dashboard environment at the same time. 	<p>If you are using IBM Installation Manager to perform the installation and configuration, first install IBM Installation Manager and setup a repository. See “Using IBM Installation Manager to install or update components from a local repository” on page 300 or “Using IBM Installation Manager to install or update components from a network repository” on page 302.</p> <p>If you are using the Tivoli Monitoring installation launchpad, see “Supported browser versions” on page 90.</p> <p>To perform the installation of the dashboard application, see “Installing the IBM Infrastructure Management Dashboards for Servers using the graphical user interface” on page 308.</p> <p>If you want to install the Tivoli Authorization Policy Server or tivcmd CLI at the same time, see also “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318.</p>
Install, update, and configure by using IBM Installation Manager silent installation mode.	“Performing a silent installation of software packages using IBM Installation Manager” on page 855
<p>Install, update, and configure by using IBM Installation Manager console mode.</p> <p>Note: You can install IBM Infrastructure Management Dashboards for Servers and the Tivoli Authorization Policy Server at the same time.</p>	<p>Install IBM Installation Manager and setup a repository using the instructions in one of the following sections: “Using IBM Installation Manager to install or update components from a local repository” on page 300 or “Using IBM Installation Manager to install or update components from a network repository” on page 302.</p> <p>To perform the installation of the dashboard application, see “Installing or updating the IBM Infrastructure Management Dashboards for Servers using console mode” on page 311.</p>
Configure the IBM Infrastructure Management Dashboards for Servers package after installation or update.	“Configuring the IBM Infrastructure Management Dashboards for Servers package after installation” on page 314
Update the IBM Infrastructure Management Dashboards for Servers package using a graphical user interface	“Updating the IBM Infrastructure Management Dashboards for Servers using the graphical user interface” on page 315
Rollback the IBM Infrastructure Management Dashboards for Servers to a previously installed version.	“Rollback of the IBM Infrastructure Management Dashboards for Servers” on page 317
Verify that the dashboard data provider is enabled in the Tivoli Enterprise Portal Server.	“Verifying the dashboard data provider is enabled” on page 427
Complete the setup and configuration of the dashboard environment.	<i>Preparing your dashboard environment in the IBM Tivoli Monitoring Administrator’s Guide</i>
Uninstall the IBM Infrastructure Management Dashboards for Servers package.	“Uninstalling IBM Tivoli Monitoring components for a dashboard environment” on page 944

Installing the IBM Infrastructure Management Dashboards for Servers using the graphical user interface

Use the IBM Tivoli Monitoring Installation Launchpad or IBM Installation Manager with a local or network repository to install the IBM Infrastructure Management Dashboards for Servers using a graphical user interface. Using both of these methods, you can also install the Tivoli Authorization Policy Server and tivcmd Command Line Interface for Authorization Server at the same time.

Before you begin

IBM Dashboard Application Services Hub must be installed before you can install Infrastructure Management Dashboards for Servers.

If IBM Dashboard Application Services Hub is not already installed, you have two options:

1. Use the Jazz for Service Management launchpad application to install IBM Dashboard Application Services Hub and its prerequisites. Then, use the IBM Tivoli Monitoring installation launchpad to install the IBM Tivoli Monitoring components for a dashboard environment. This is the recommended approach if you are not familiar with IBM Installation Manager.
2. Ensure that IBM Installation Manager is installed and then set up a local or network repository that contains the Jazz for Service Management components. Also set up a repository with the IBM Tivoli Monitoring components for a dashboard environment. Then, use IBM Installation Manager to install IBM Dashboard Application Services Hub, its prerequisites, and the IBM Tivoli Monitoring components at the same time. IBM Installation Manager ensures that components are installed in the correct order.

If Dashboard Application Services Hub is already installed, you have two options for installing Infrastructure Management Dashboards for Servers:

1. Use the IBM Tivoli Monitoring installation launchpad. This is the recommended approach if you are not familiar with IBM Installation Manager.
2. Set up a local or network repository with the Tivoli Monitoring components for a dashboard environment and then use IBM Installation Manager to perform the installation.

If you are using an IBM Installation Manager repository to perform the installation, you must ensure that IBM Installation Manager is installed and then set up a local or network repository with the packages to be installed. For more information, see “Using IBM Installation Manager to install or update components from a local repository” on page 300 or “Using IBM Installation Manager to install or update components from a network repository” on page 302.

Note: If you select the *Configuration* feature in this procedure, the installer will restart the IBM WebSphere Application Server as part of the installation process.

About this task

Complete the following procedure to install and configure the infrastructure management dashboards on the computer where the Dashboard Application Services Hub is installed. (Or where the Dashboard Application Services Hub will be installed at the same time as the infrastructure management dashboards.)

Procedure

1. If you are using the IBM Tivoli Monitoring Installation Launchpad, perform the following steps to start the installation:
 - a. Log in to the system as the same user who installed the IBM Dashboard Application Services Hub.

- b. Execute the launchpad command or script for your operating system. The launchpad commands and scripts are located in the root directory of the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image.
 - **Windows** 32 bit system launchpad.exe
 - **Windows** 64 bit system launchpad64.exe
 - **Linux** **UNIX** launchpad.sh
- c. On the main launchpad panel, select one of the following links under the Product Overview category:
 - *Install as administrative user*: Select this link if you installed IBM Dashboard Application Services Hub as an administrative user on Windows or as root on Linux/UNIX.
 - *Install as non-administrative user*: Select this link if you did not install IBM Dashboard Application Services Hub as an administrative user on Windows or as root on Linux/UNIX.
- d. In the Installation Manager main window, click **Install** to see the available packages in the repositories that are configured for IBM Installation Manager.

The IBM Installation Manager window opens with several choices for managing your IBM software installations.

2. If you are using IBM Installation Manager, perform these steps to start the installation:
 - a. Log in as the user who installed IBM Dashboard Application Services Hub if it is already installed.
 - b. Start the IBM Installation Manager:
 - **Windows** Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager.
 - **Linux** **UNIX** Execute the IBMIM binary under /opt/IBM/InstallationManager/eclipse.

The IBM Installation Manager window opens with several choices for managing your IBM software installations.
 - c. In the Installation Manager main window, click **Install** to see the available packages in the repositories that are configured for IBM Installation Manager.
3. Select the **Infrastructure Management Dashboards for Servers** package and click **Next**.

Note:

- If you are using the IBM Tivoli Monitoring Installation Launchpad, the **Infrastructure Management Dashboards for Servers** package is preselected.
 - If you are using IBM Installation Manager to install IBM Dashboard Application Services Hub at the same time, you must also select its package and its prerequisite packages. See the Jazz for Services Management Installation Guide for details when you are asked to provide information specific to these packages (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html).
 - You can also select the Tivoli Authorization Policy Server package and tivcmd Command Line Interface for Authorization Policy packages, if they are also being installed on the same system. If you select the Tivoli Authorization Policy Server package, see “Installing and configuring the Tivoli Authorization Policy Server using the graphical user interface” on page 320 for details when you are asked to provide information specific to these packages.
4. If you are installing Jazz for Service Management, accept the license agreement for each package and click **Next**.
 5. If you are using IBM Installation Manager to perform the installation, specify the location of the Shared Resources Directory and click **Next**.

The first time that you install a package on a computer you are asked to specify a shared resources directory. The shared resources directory is where IBM Installation Manager stores installation

artifacts that can be used by one or more package groups. Choose a directory on your largest drive. You cannot change the shared resources directory location until you uninstall all packages.

6. On the panel for configuring package groups, you cannot specify an installation directory for the **Infrastructure Management Dashboards for Servers** package because it is installed using the package group definition for IBM Dashboard Application Services Hub: **Core services in Jazz for Service Management**.

Note:

- If you are using IBM Installation Manager to install IBM Dashboard Application Services Hub at the same time, you should confirm the package group location for its package group and its prerequisite package groups.
- If you are also installing the tivcmd Command Line Interface for Authorization Policy package on the computer, you should select its package directory and either use the default installation directory location or enter a custom installation directory location. On a 64 bit machine, the 64 bit architecture of the tivcmd Command Line Interface for Authorization Policy is installed by default. You can select to switch to the 32 bit architecture. However, on zSeries systems with Linux, you can only install the 32 bit architecture of the CLI.

Important: The tivcmd CLI cannot be installed under the same installation directory as Jazz for Service Management or the monitoring server, portal server, portal client, monitoring agents, and tacmd CLI components of IBM Tivoli Monitoring.

7. If you are using IBM Installation Manager to perform the installation for Jazz for Service Management, you might be asked to specify the translation packages to install. After making your selection, click **Next**.

Note: The language files for all supported languages are always installed with the IBM Infrastructure Management Dashboards for Servers, Tivoli Authorization Policy Server, and tivcmd Command Line Interface for Authorization Policy packages.

8. Select the **Infrastructure Management Dashboards for Servers** feature and click **Next**.

Note:

- If you are using the IBM Tivoli Monitoring Installation Launchpad, this feature is preselected.
 - It is recommended to select the **Infrastructure Management Dashboards for Servers** Installation and Configuration features together (which is the default behavior). However, you can clear the Configuration feature if you did not select the Configuration feature when IBM Dashboard Application Services Hub was installed. Or likewise, if you want to review the deployment scripts for the Infrastructure Management Dashboards for Servers before they are used to deploy the infrastructure management dashboards application into IBM Dashboard Application Services Hub. If you select the Installation feature on its own, only the dashboard application binaries and install scripts are laid down in the installation directory (on Windows this location is usually C:\Program Files\IBM\JazzSM\installedDashboards\com.ibm.tivoli.itmosd and on Linux/UNIX this location is usually /opt/IBM/JazzSM/installedDashboards/com.ibm.tivoli.itmosd). If you select the Configuration feature on its own, the Installation feature is selected as well.
 - If you are using IBM Installation Manager to install IBM Dashboard Application Services Hub at the same time, you must also select the IBM Dashboard Application Services Hub features and its prerequisite features.
 - You can also select the Tivoli Authorization Policy Server feature and tivcmd Command Line Interface for Authorization Policy feature if you want to install those features on the same system.
9. On the next window, specify the configuration parameters for each package where you selected the Configuration feature. Select each package on the left to see the set of configuration parameters. Specify the configuration parameters and click **Next**.

Note:

- The **Next** button is not enabled until all required configuration parameters are completed.
 - If you are using IBM Installation Manager to install IBM Dashboard Application Services Hub at the same time as the infrastructure management dashboards, see the Jazz for Service Management Installation Guide at http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html for details when you are asked to provide information specific to these packages.
 - If IBM Dashboard Application Services Hub is already installed, you must specify an IBM Dashboard Application Services Hub administrative user name and password, such as `smadmin` and its password for the Common Configurations for Core services in Jazz for Service Management feature, and click **Validate**. These credentials are used to deploy the Infrastructure Management Dashboards for Servers application into the WebSphere Application Server for IBM Dashboard Application Services Hub.
 - If you are also installing the Tivoli Authorization Policy Server and `tivcmd` Command Line Interface for Authorization Policies at the same time, see step 9 in “Installing and configuring the Tivoli Authorization Policy Server using the graphical user interface” on page 320 for more details on the configuration parameters for these components.
10. Click **Next** to review the target location and a summary of packages to install. This is your opportunity to click **Back** and make any changes before proceeding.
 11. Click **Install** to begin installing. After installation is complete, the packages are shown in a panel with a message confirming that they have been installed. At this point, you can review the log files.
 12. Click **Finish** to finalize the installation.

What to do next

If you selected the Installation and Configuration features of the IBM Infrastructure Management Dashboards for Servers, perform the following tasks:

1. Verify the dashboard data provider is enabled in the portal server configuration. See “Verifying the dashboard data provider is enabled” on page 427.
2. Then, see *Preparing your dashboard environment* in the *IBM Tivoli Monitoring Administrator's Guide* for additional configuration steps that are required to set up authentication and authorization for dashboard users. This topic also contains information about creating a data provider connection from IBM Dashboard Application Services Hub to the IBM Tivoli Monitoring dashboard data provider component of the portal server.
3. For more information about launching and using IBM Infrastructure Management Dashboard for Servers, see the *Windows OS Agent User's Guide*, the *Linux OS Agent User's Guide*, or the *UNIX OS Agent User's Guide*.

If you only selected the Installation feature of the IBM Infrastructure Management Dashboards for Servers, review the dashboard application's deployment scripts and then use the **Modify** operation of IBM Installation Manager to deploy the dashboard application into IBM Dashboard Application Services Hub. See “Configuring the IBM Infrastructure Management Dashboards for Servers package after installation” on page 314.

Installing or updating the IBM Infrastructure Management Dashboards for Servers using console mode

You can use IBM Installation Manager in console mode to install or update packages from a command-line interface.

Before you begin

Before installing or updating the IBM Infrastructure Management Dashboards for Servers, install or update the prerequisite software that is described in “Required software and memory requirements for a dashboard environment” on page 91.




Set up a local or network repository that contains the IBM Tivoli Monitoring components for a dashboard environment. See “Using IBM Installation Manager to install or update components from a local repository” on page 300 and “Using IBM Installation Manager to install or update components from a network repository” on page 302.

Note: If you select the *Configuration* feature in this procedure, the installer will restart the IBM WebSphere Application Server as part of the installation or update process.

About this task

Perform the following steps to run the IBM Installation Manager in console mode to install or update IBM Infrastructure Management Dashboards for Servers on the computer where IBM Dashboard Application Services Hub is installed.

Procedure

1. Log in as the user who installed IBM Dashboard Application Services Hub.
2. Start the IBM Installation Manager in console mode from the directory where it was installed. For example:
 -  `C:\Program Files\IBM\Installation Manager\eclipse\tools\imcl.exe -c`
 -   `/opt/IBM/InstallationManager/eclipse/tools/imcl -c`
3. The following menu is displayed:

```
Select:
  1. Install - Install software packages
  2. Update - Find and install updates and fixes to installed software packages
  3. Modify - Change installed software packages
  4. Roll Back - Revert to an earlier version of installed software packages
  5. Uninstall - Remove installed software packages

Other Options:
  L. View Logs
  S. View Installation History
  V. View Installed Packages
  -----
  P. Preferences
  -----
  E. Export Data for Problem Analysis
  A. About IBM Installation Manager
  -----
  X. Exit Installation Manager
```

Select **1** to install software packages. Select **2** to update software packages.

4. When you are presented with the list of software packages, select the option for the **IBM Infrastructure Management Dashboards for Servers** package.

Note:

- a. You can also select the Tivoli Authorization Policy Server if you want to install it on the same system. If you select the Tivoli Authorization Policy Server package, see “Installing and configuring the Tivoli Authorization Policy Server using console mode” on page 325 for details when you are asked to provide information specific to this package.
- b. You should not select the tivcmd Command Line Interface package since it cannot be installed at the same time as the Authorization Policy Server or infrastructure management dashboards by using console mode.
- c. When you update packages, the IBM Installation Manager only displays the packages with an updated version in the repository. For example, even though you might have Infrastructure Management Dashboards for Servers and the Authorization Policy Server installed, you might only see the updated dashboard package if the Authorization Policy Server was not updated in the latest release or fix pack.

5. From the next list of options, select 1 to install or update the current version of the software package. Option 2 displays all available versions of the package.
6. IBM Infrastructure Management Dashboards for Servers is selected by default. Select N (Next) to continue.
7. Select the IBM Infrastructure Management Dashboards for Servers features to install or update. By default, both *Installation* and *Configuration* are selected. Select N (Next) to continue.

Note: It is recommended to select the Infrastructure Management Dashboards for Servers *Installation* and *Configuration* features together (which is the default behavior). However, you cannot select the infrastructure management dashboards Configuration feature if you did not select the Configuration feature when IBM Dashboard Application Services Hub was installed. You also should not select the Configuration feature if you want to review the deployment scripts before they are used to deploy the infrastructure management dashboards application into IBM Dashboard Application Services Hub. If you select the Installation feature on its own, only the dashboard application binaries and install scripts are laid down in the installation directory (on Windows this location is usually C:\Program Files\IBM\JazzSM\installedDashboards\com.ibm.tivoli.itmosd and on Linux/UNIX this location is usually /opt/IBM/JazzSM/installedDashboards/com.ibm.tivoli.itmosd). If you select the Configuration feature on its own, the Installation feature is selected as well.

8. Enter your WebSphere administrator-level user credentials (user name and password) to install or update the Infrastructure Management Dashboards for Servers application into IBM Dashboard Application Services Hub. The installer validates the entered credentials.

Note: If you are installing or updating other packages, you must supply configuration information for those packages. See step 10 in “Installing and configuring the Tivoli Authorization Policy Server using console mode” on page 325 for information about configuring the Tivoli Authorization Policy Server feature.

9. When validation is completed select N (Next) to continue.
10. Review the installation summary window and select I to proceed with the installation or update.

What to do next

If you are installing IBM Infrastructure Management Dashboards for Servers for the first time and you selected the Installation and Configuration features, perform the following additional tasks:

1. Verify the dashboard data provider is enabled in the portal server configuration. See “Verifying the dashboard data provider is enabled” on page 427.
2. Then, see *Preparing your dashboard environment* in the *IBM Tivoli Monitoring Administrator's Guide* for additional configuration steps that are required to set up authentication and authorization for dashboard users. This topic also contains information about creating a data provider connection from IBM Dashboard Application Services Hub to the IBM Tivoli Monitoring dashboard data provider component of the portal server.
3. For more information about launching and using IBM Infrastructure Management Dashboard for Servers, see the *Windows OS Agent User's Guide*, the *Linux OS Agent User's Guide*, or the *UNIX OS Agent User's Guide*.

If you only selected the Installation feature of the IBM Infrastructure Management Dashboards for Servers, review the dashboard application's deployment scripts and then use the **Modify** operation of IBM Installation Manager to deploy the dashboard application into IBM Dashboard Application Services Hub. See “Configuring the IBM Infrastructure Management Dashboards for Servers package after installation” on page 314.

Configuring the IBM Infrastructure Management Dashboards for Servers package after installation




If you chose the Installation feature alone for the IBM Infrastructure Management Dashboards for Servers at installation or update time, use the **Modify** option in IBM Installation Manager for configuration after you have reviewed the deployment scripts of the dashboard application. When you configure the dashboard application, it is deployed to IBM Dashboard Application Services Hub.

Before you begin

Ensure that the configuration feature for IBM Dashboard Application Services Hub is already selected. IBM Infrastructure Management Dashboard for Servers is an extension offering of the Dashboard Application Services Hub. As a result, an error message is displayed if you attempt to select the infrastructure management dashboard configuration feature when Dashboard Application Servers Hub is not configured.

Note: The installer will restart the IBM WebSphere Application Server as part of the installation or update process.

Procedure

1. Log in as the user who installed IBM Dashboard Application Services Hub.
2. On the system where the Dashboard Application Services Hub and Infrastructure Management Dashboards for Servers are installed, start the IBM Installation Manager:
 -  Double-click the IBMIM.exe file in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager.
 -   Execute the IBMIM binary under /opt/IBM/InstallationManager/eclipse.

The IBM Installation Manager window opens with several choices for managing your IBM software installations.

3. Click **Modify** to see the available features for configuring.
4. Select **Core services in Jazz for Service Management**. This option contains the **IBM Infrastructure Management Dashboards for Servers** package. Click **Next** to continue.
5. Expand the **IBM Infrastructure Management Dashboards for Servers** node and select the **Configuration** check box. Click **Next** to continue.
6. Click **Next**. The common configuration information for the Dashboard Application Services Hub is shown in the fields.
7. Specify an IBM Dashboard Application Services Hub administrative username and password, such as smadmin and its password for the **Common Configurations for Core services in Jazz for Service Management** feature, and click **Validate**. These credentials are used to deploy the **Infrastructure Management Dashboards for Servers** application into the WebSphere Application Server for IBM Dashboard Application Services Hub.

After your credentials are successfully validated, the **Next** button is enabled.
8. Click **Next** to see and review the summary information. This is your opportunity to click **Back** and make any changes before proceeding.
9. Click **Modify** to deploy the infrastructure management dashboards to the Dashboard Application Services Hub. After configuration is complete, the packages are shown in a panel with a message confirming the deployment.
10. Click **Finish** to finalize the configuration.

What to do next

If you are installing IBM Infrastructure Management Dashboards for Servers for the first time, perform the following additional tasks:

1. Verify the dashboard data provider is enabled in the portal server configuration. See “Verifying the dashboard data provider is enabled” on page 427.
2. Then, see *Preparing your dashboard environment* in the *IBM Tivoli Monitoring Administrator's Guide* for additional configuration steps that are required to set up authentication and authorization for dashboard users. This topic also contains information about creating a data provider connection from IBM Dashboard Application Services Hub to the IBM Tivoli Monitoring dashboard data provider component of the portal server.
3. For more information about launching and using IBM Infrastructure Management Dashboard for Servers, see the *Windows OS Agent User's Guide*, the *Linux OS Agent User's Guide*, or the *UNIX OS Agent User's Guide*.

Updating the IBM Infrastructure Management Dashboards for Servers using the graphical user interface

Use the IBM Tivoli Monitoring Installation Launchpad or IBM Installation Manager with a local or network repository to update an existing installation of IBM Infrastructure Management Dashboards for Servers to a later version such as a Fix Pack version.

Before you begin

See “Required software and memory requirements for a dashboard environment” on page 91 to determine if you need to update any of the prerequisite software before updating IBM Infrastructure Management Dashboards for Servers.





If you are not using the IBM Tivoli Monitoring Installation Launchpad, set up a local or network repository that contains the IBM Tivoli Monitoring components for a dashboard environment. See “Using IBM Installation Manager to install or update components from a local repository” on page 300 and “Using IBM Installation Manager to install or update components from a network repository” on page 302.

Note: If you select the *Configuration* feature in this procedure, the installer will restart the IBM WebSphere Application Server as part of the update process.

About this task

Complete the following steps to update the IBM Infrastructure Management Dashboards for Servers using a graphical user interface.

Procedure

1. If you are using the IBM Tivoli Monitoring Installation Launchpad, perform the following steps to start the update:
 - a. Log in to the system as the same user who installed the IBM Dashboard Application Services Hub.
 - b. Execute the launchpad command or script for your operating system. The launchpad commands and scripts are located in the root directory of the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image.
 -  32 bit system `launchpad.exe`
 -  64 bit system `launchpad64.exe`
 -   `launchpad.sh`
 - c. On the main launchpad panel, select one of the following links under the Product Overview category:

- *Install as administrative user:* Select this link if you installed IBM Dashboard Application Services Hub as an administrative user on Windows or as root on Linux/UNIX.
- *Install as non-administrative user:* Select this link if you did not install IBM Dashboard Application Services Hub as an administrative user on Windows or as root on Linux/UNIX.

The IBM Installation Manager window opens with several choices for managing your IBM software installations.

2. If you are using IBM Installation Manager, perform these steps to start the update:
 - a. Log in as the user who installed IBM Dashboard Application Services Hub.
 - b. Start the IBM Installation Manager:
 - **Windows** Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager.
 - **Linux** **UNIX** Execute the IBMIM binary under /opt/IBM/InstallationManager/eclipse.

The IBM Installation Manager window opens with several choices for managing your IBM software installations.

3. In the Installation Manager main window, click **Update** to see the available packages in the repositories that are configured for IBM Installation Manager.
4. Select the **Core services in Jazz for Service Management** package group to find updates. Click **Next** to continue.
5. Installation Manager displays the version of the updated Infrastructure Management Dashboards for Servers package. Click **Next** to continue.

Note: When you update packages, the IBM Installation Manager only displays the packages with an updated version in the repository. For example, even though you might have Infrastructure Management Dashboards for Servers and the Authorization Policy Server installed, you might only see the updated dashboard package if the Authorization Policy Server was not updated in the latest release or fix pack.

6. Select the Installation and Configuration features of the Infrastructure Management Dashboards for Servers package and click **Next**.

Note:

- It is recommended to select the **Infrastructure Management Dashboards for Servers** Installation and Configuration features together (which is the default behavior). Clear the Configuration feature if you want to review the deployment scripts for the Infrastructure Management Dashboards for Servers before they are used to deploy the infrastructure management dashboards application into IBM Dashboard Application Services Hub. If you select the Installation feature on its own, only the dashboard application binaries and install scripts are laid down in the installation directory (on Windows this location is usually C:\Program Files\IBM\JazzSM\installedDashboards\com.ibm.tivoli.itmosd and on Linux/UNIX this location is usually /opt/IBM/JazzSM/installedDashboards/com.ibm.tivoli.itmosd). If you select the Configuration feature on its own, the Installation feature is selected as well.
7. On the next window, specify the configuration parameters for each package where you selected the Configuration feature. Select each package on the left to see the set of configuration parameters. Specify the configuration parameters and click **Next**.

Note:

- The **Next** button is not enabled until all required configuration parameters are completed.
- You must specify an IBM Dashboard Application Services Hub administrative user name and password, such as smadmin and its password for the Common Configurations for Core services in Jazz for Service Management feature, and click **Validate**. These credentials are used to deploy the

Infrastructure Management Dashboards for Servers application into the WebSphere Application Server for IBM Dashboard Application Services Hub.

8. Click **Next** to review the target location and a summary of packages to update. This is your opportunity to click **Back** and make any changes before proceeding.
9. Click **Update** to update the IBM Infrastructure Management Dashboards for Servers.

What to do next

If you only selected the Installation feature of the IBM Infrastructure Management Dashboards for Servers, review the dashboard application's deployment scripts and then use the **Modify** operation of IBM Installation Manager to deploy the updated dashboard application into IBM Dashboard Application Services Hub. See “Configuring the IBM Infrastructure Management Dashboards for Servers package after installation” on page 314.

For more information about launching and using IBM Infrastructure Management Dashboard for Servers, see the *Windows OS Agent User's Guide*, the *Linux OS Agent User's Guide*, or the *UNIX OS Agent User's Guide*.

Rollback of the IBM Infrastructure Management Dashboards for Servers

You can use the rollback feature of IBM Installation Manager to revert an updated package back to a previous version. To complete the rollback, the Installation Manager requires information from the earlier version of the package, if the package supports rollback.

Before you begin

When Installation Manager rolls a package back to a previous version, it uninstalls the current version of some files and reinstalls the previous version of those files.

By default, the files that are required to complete the rollback are saved on your computer. On the **Files for Rollback** preference page, you can change the default action and remove any previously stored files. For more information, see “Saving package information for rollback” on page 305.

Note: The installer will restart the IBM WebSphere Application Server as part of the rollback process.

Restriction: IBM Installation Manager does not store your Jazz for Service Management profile location and password in the response file that is recorded for rollback. You must manually update the recorded response file with the missing profile location and password. The following entry is an example of the missing section:

```
<profile id='Core services in Jazz for Service Management' installLocation='C:\Program Files\IBM\JazzSM'>
  <data key='user.WAS_PASSWORD' value='M6NQ+7Gb+GuoNQNg8N+Lkw==' />
</profile>
```

Where:

installLocation='C:\Program Files\IBM\JazzSM'

Is the customized installation location.

value='M6NQ+7Gb+GuoNQNg8N+Lkw=='

Is the encrypted password for the user.

For more information, see “Recording a response file” on page 858.

About this task

Complete the following steps to rollback the IBM Infrastructure Management Dashboards for Servers if IBM Installation Manager was configured to save the files for the previous version of the package:

Procedure

1. On the main Installation Manager window, click **Rollback** to start the Rollback wizard.
2. Select the **Core services in Jazz for Service Management** package group and click **Next**.
3. Select the version of the package that you want to rollback to and click **Next**.
4. Review the summary information and click **Rollback** to rollback the package.
5. Follow the steps that are displayed in the wizard. These are the same steps as “Updating the IBM Infrastructure Management Dashboards for Servers using the graphical user interface” on page 315.

Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy

Authorization policies provide you with role-based access control capabilities for monitoring dashboards in IBM Dashboard Application Services Hub. You can restrict the access that dashboard users have to managed system groups and to individual agent managed systems. The Tivoli Authorization Policy Server and tivcmd Command-Line Interface (CLI) for Authorization Policy are required framework components for the IBM Tivoli Monitoring authorization policy feature.

For more information and advantages of using authorization policies over Tivoli Enterprise Portal user permissions to control dashboard user access to managed system groups and managed systems, see “Authorization and authentication” on page 139.

The Authorization Policy Server is installed with IBM Dashboard Application Services Hub.

Note: The IBM Dashboard Application Services Hub must be restarted after the Authorization Policy Server is configured.

If you plan to use authorization policies, normally you install the Authorization Policy Server with each IBM Dashboard Application Services Hub where you also install the IBM Infrastructure Management Dashboards for Servers. Or where you install other dashboard monitoring applications such as IBM Infrastructure Management Dashboards for VMware, or where you plan to create custom monitoring dashboards that use the IBM Tivoli Monitoring dashboard data provider. However, the following exceptions apply:

1. If you set up multiple Dashboard Application Services Hub Servers for load balancing with an IBM HTTP Server, you cannot install the Authorization Policy Server into this environment since it does not support load balancing with this release.
2. If you have multiple IBM Tivoli Monitoring environments (also called domains) and you can share a set of common authorization policies between the monitoring domains, you should not install the Authorization Policy Server with each IBM Dashboard Application Services Hub. For more information, see “Multi-hub environments” on page 63.

The Tivoli Authorization Policy Server package has two features that you can select during installation, **Installation** and **Configuration**. When you choose the Installation feature, the Authorization Policy Server files are copied from the installation media or repository to the file system. When you choose the Configuration feature, the Authorization Policy Server application is deployed into IBM Dashboard Application Services Hub using the scripts that were installed by the Installation feature. Normally you should select both features during installation. However, if you want to view the deployment scripts before they are executed, you can choose the Installation feature alone during installation. After

reviewing the deployment scripts, you then use the Modify operation of IBM Installation Manager to select the Configuration feature and deploy the Authorization Policy Server application into IBM Dashboard Application Services Hub.

Note: Although you could use the WebSphere Application Server wsadmin command-line interface to execute the Authorization Policy Server deployment scripts, this method is not supported. IBM Installation Manager must be used to track which features are installed and configured.

The tivcmd CLI should be installed on the computers where administrators create and work with authorization policies. The tivcmd CLI connects to an IBM Dashboard Application Services Hub where the Tivoli Authorization Policy Server package is installed by using HTTP or HTTPS. Therefore, it can either be installed on the same computer as Authorization Policy Server or on other computers that administrators use.

Authorization policy checking is disabled by default at the Tivoli Enterprise Portal Server. The installation of the Tivoli Authorization Policy Server and tivcmd CLI have no effect on IBM Tivoli Monitoring runtime behavior unless you enable authorization policy checking during configuration or reconfiguration of the Tivoli Enterprise Portal Server.

After successful installation of the Tivoli Authorization Policy Server and tivcmd Command Line Interface for Authorization Policy packages, you can execute various tivcmd CLI commands as required to create roles, grant permissions, exclude permissions, and so on. The *IBM Tivoli Monitoring Administrator's Guide* also provides examples of creating authorization policies for common scenarios in the *Using role-based authorization policies* chapter. The *Preparing your dashboard environment* chapter contains steps for configuring a dashboard environment to use authorization policies. For a complete list of tivcmd CLI commands, see the *IBM Tivoli Monitoring Command Reference*.

For a complete list of prerequisite criteria, see “Required software and memory requirements for a dashboard environment” on page 91.

See the following table for information about the tasks that you perform to install and configure the Authorization Policy Server and tivcmd CLI:

Table 59. Installation and configuration topics

Task	Where to find detailed information
<p>Install and configure using the IBM Tivoli Monitoring installation launchpad or the IBM Installation Manager graphical user interface.</p> <p>Note:</p> <ol style="list-style-type: none"> 1. The Tivoli Monitoring installation launchpad is the recommended approach if you are not familiar with IBM Installation Manager and you have a supported browser on the computer system where you are installing. 2. You can install one or more of the IBM Tivoli Monitoring packages for a dashboard environment at the same time. 	<p>If you are using IBM Installation Manager to perform the installation and configuration, first install IBM Installation Manager and set up a repository. See “Using IBM Installation Manager to install or update components from a local repository” on page 300 and “Using IBM Installation Manager to install or update components from a network repository” on page 302.</p> <p>If you are using the Tivoli Monitoring installation launchpad, see “Supported browser versions” on page 90.</p> <p>To install the Authorization Policy Server or tivcmd CLI, see “Installing and configuring the Tivoli Authorization Policy Server using the graphical user interface” on page 320 and “Installing the tivcmd Command Line Interface for Authorization Policy using the graphical user interface” on page 330.</p> <p>If you plan to install the infrastructure management dashboards, also see “Installing and configuring the IBM Infrastructure Management Dashboards for Servers” on page 306</p>

Table 59. Installation and configuration topics (continued)

Task	Where to find detailed information
Install and configure using IBM Installation Manager console mode. Note: You can install the Tivoli Authorization Policy Server and IBM Infrastructure Management Dashboards for Servers at the same time but the tivcmd CLI package must be installed separately.	Set up a repository by using the instructions in one of the following sections: “Using IBM Installation Manager to install or update components from a local repository” on page 300 or “Using IBM Installation Manager to install or update components from a network repository” on page 302. To install the Authorization Policy Server, see “Installing and configuring the Tivoli Authorization Policy Server using console mode” on page 325. To install the tivcmd command line interface, see “Installing the tivcmd Command Line Interface for Authorization Policy using console mode” on page 332.
Install and configure using IBM Installation Manager silent installation mode.	“Performing a silent installation of software packages using IBM Installation Manager” on page 855
Configure the Tivoli Authorization Policy Server feature after installation.	“Configuring the Tivoli Authorization Policy Server feature after installation” on page 328
Verify the Tivoli Authorization Policy Server installation.	“Verifying the Tivoli Authorization Policy Server installation” on page 329
Create authorization policies and complete the set up and configuration of the dashboard environment to use authorization policies. Note: The user credentials that are specified during the installation and configuration of the Authorization Policy Server are assigned to the PolicyAdministrator role. You must use these credentials with the tivcmd Command Line Interface to log in to the Authorization Policy Server and assign other administrators permission to create and work with authorization policies.	<i>Using role-based authorization policies and Preparing your dashboard environment in the IBM Tivoli Monitoring Administrator’s Guide.</i>
Use IBM Installation Manager to save the Authorization Policy Server package information that is needed to perform a rollback in the future.	“Saving package information for rollback” on page 305
Upgrade the Tivoli Authorization Policy Server.	“Updating the Tivoli Authorization Policy Server using the graphical user interface” on page 334
Rollback the Tivoli Authorization Policy Server to a previously installed version.	“Rollback of the Tivoli Authorization Policy Server” on page 336
Uninstall procedures.	“Uninstalling IBM Tivoli Monitoring components for a dashboard environment” on page 944

Installing and configuring the Tivoli Authorization Policy Server using the graphical user interface

Use the IBM Tivoli Monitoring Installation Launchpad or IBM Installation Manager with a local or network repository to install the Tivoli Authorization Policy Server using a graphical user interface. Using both of these methods, you can also install the IBM Infrastructure Management Dashboards for Servers and tivcmd Command Line Interface for Authorization at the same time. The Installation Launchpad is the recommended approach if you do not have experience using IBM Installation Manager.

Before you begin

IBM Dashboard Application Services Hub must be installed before you can install Authorization Policy Server.

If IBM Dashboard Application Services Hub is not already installed, you have two options:

1. Use the Jazz for Service Management launchpad application to install IBM Dashboard Application Services Hub and its prerequisites. Then, use the IBM Tivoli Monitoring installation launchpad to install the IBM Tivoli Monitoring components for a dashboard environment. This is the recommended approach if you are not familiar with IBM Installation Manager.
2. Ensure that IBM Installation Manager is installed and then set up a local or network repository that contains the Jazz for Service Management components. Also set up a repository with the IBM Tivoli Monitoring components for a dashboard environment. Then use IBM Installation Manager to install IBM Dashboard Application Services Hub, its prerequisites, and the IBM Tivoli Monitoring components at the same time. IBM Installation Manager ensures the components are installed in the correct order.

If Dashboard Application Services Hub is already installed, you have two options for installing the Tivoli Authorization Policy Server:

1. Use the IBM Tivoli Monitoring installation launchpad. This is the recommended approach if you are not familiar with IBM Installation Manager.
2. Set up a local or network repository with the Tivoli Monitoring components for a dashboard environment and then use IBM Installation Manager to perform the installation.

If you are using IBM Installation Manager repository to perform the installation, you must ensure that it is installed and then set up a local or network repository with the packages to be installed. For more information, see “Using IBM Installation Manager to install or update components from a local repository” on page 300 or “Using IBM Installation Manager to install or update components from a network repository” on page 302.

About this task

Complete the following steps to install and configure the Tivoli Authorization Policy Server on the computer where the Dashboard Application Services Hub is installed. (Or where the Dashboard Application Services Hub will be installed at the same time as the authorization policy server.)

Procedure

1. If you are using the IBM Tivoli Monitoring Installation Launchpad, perform these steps to start the installation:
 - a. Log in into the system as the same user who installed the IBM Dashboard Application Services Hub.
 - b. Execute the launchpad command or script for your operating system. The launchpad commands and scripts are located in the root directory of the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components DVD* or DVD image.
 - **Windows** 32 bit system `launchpad.exe`
 - **Windows** 64 bit system `launchpad64.exe`
 - **Linux** **UNIX** `launchpad.sh`
 - c. On the main launchpad panel, select one of the following links under the Product Overview category:
 - *Install as administrative user*: Select this link if you installed IBM Dashboard Application Services Hub as an administrative user on Windows or as root on Linux/UNIX.
 - *Install as non-administrative user*: Select this link if you did not install IBM Dashboard Application Services Hub as an administrative user on Windows or as root on Linux/UNIX.
 - d. In the Installation Manager main window, click **Install** to display the available packages.

The IBM Installation Manager window opens with several choices for managing your IBM software installations.

2. If you are using IBM Installation Manager, perform the following steps to start the installation:

- a. Log in as the user who installed IBM Dashboard Application Services Hub if it is already installed.
- b. Start the IBM Installation Manager:
 - **Windows** Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager\eclipse.
 - **Linux** **UNIX** Execute the IBMIM binary under /opt/IBM/InstallationManager/eclipse.

The IBM Installation Manager window opens with several choices for managing your IBM software installations.
- c. In the Installation Manager main window, click **Install** to see the available packages in the repositories that are configured for IBM Installation Manager.
3. Select the **Tivoli Authorization Policy Server** package and click **Next**.

Note:

- If you are using the IBM Tivoli Monitoring Installation Launchpad, the **Tivoli Authorization Policy Server** package is preselected.
 - If you are using IBM Installation Manager to install IBM Dashboard Application Services Hub at the same time, you must also select its package and its prerequisite packages. See the Jazz for Services Management Installation Guide for details when you are asked to provide information specific to these packages (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html).
 - You can also select the IBM Infrastructure Management Dashboards for Servers package and tivcmd Command Line Interface for Authorization Policy packages if they are also being installed on the same system.
4. If you are installing Jazz for Service Management, accept the license agreement for each package and click **Next**.
 5. If you are using IBM Installation Manager to perform the installation, specify the location of the Shared Resources Directory and click **Next**.
 The first time that you install a package on a computer you are asked to specify a shared resources directory. The shared resources directory is where IBM Installation Manager stores installation artifacts that can be used by one or more package groups. Choose a directory on your largest drive. You cannot change the shared resources directory location until you uninstall all packages.
 6. On the panel for configuring package groups, you cannot specify an installation directory for the **Tivoli Authorization Policy Server** package because it is installed using the package group definition for IBM Dashboard Application Services Hub: **Core services in Jazz for Service Management**.

Note:

- If you are using IBM Installation Manager to install IBM Dashboard Application Services Hub at the same time, you should confirm the package group location for its package group and its prerequisite package groups.
- If you are also installing the tivcmd Command Line Interface for Authorization Policy package on the computer, you should select its package directory and either use the default installation directory location or enter a custom installation directory location. On a 64 bit machine, the 64 bit architecture of the tivcmd Command Line Interface for Authorization Policy is installed by default. You can select to switch to the 32 bit architecture. However, on zSeries systems with Linux, you can only install the 32 bit architecture of the CLI.

Important: The tivcmd CLI cannot be installed under the same installation directory as Jazz for Service Management or the monitoring server, portal server, portal client, monitoring agents, and tacmd CLI components of IBM Tivoli Monitoring.

7. If you are using IBM Installation Manager to perform the installation for Jazz for Service Management, you might be asked to specify the translation packages to install. After making your selection, click **Next**.

Note: The language files for all supported languages are always installed with the IBM Infrastructure Management Dashboards for Servers, Tivoli Authorization Policy Server, and tivcmd Command Line Interface for Authorization Policy packages.

8. Select the **Tivoli Authorization Policy Server** feature and click **Next**.

Note:

- If you are using the IBM Tivoli Monitoring Installation Launchpad, this feature is preselected.
 - It is recommended to select the **Tivoli Authorization Policy Server** *Installation* and *Configuration* features together (which is the default behavior). However, you can clear the Configuration feature if you did not select the Configuration feature when IBM Dashboard Application Services Hub was installed. Or likewise, if you want to review the deployment scripts before they are used to deploy the authorization policy server application into IBM Dashboard Application Services Hub. If you select the Installation feature on its own, only the authorization policy server binaries and install scripts are laid down in the installation directory (on Windows this location is usually C:\Program Files\IBM\JazzSM\AuthPolicyServer and on Linux/UNIX this location is usually /opt/IBM/JazzSM/AuthPolicyServer). If you select the Configuration feature on its own, the Installation feature is selected as well.
 - If you are using IBM Installation Manager to install IBM Dashboard Application Services Hub at the same time, you must also select the IBM Dashboard Application Services Hub features and its prerequisite features.
 - You can also select the IBM Infrastructure Management Dashboards for Servers feature and tivcmd Command Line Interface for Authorization Policy feature if you want to install those features on the same system.
9. On the next window, specify the configuration parameters for each package where you selected the Configuration feature. Select each package on the left to see its set of configuration parameters. Specify the configuration parameters and click **Next**.

Note: If you are using IBM Installation Manager to install IBM Dashboard Application Services Hub at the same time as the Tivoli Authorization Policy Server, see the Jazz for Service Management Installation Guide at http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/install/psc_c_install.html for details when you are asked to provide information specific to these packages.

The following configuration parameters apply to the Tivoli Authorization Policy Server configuration feature:

- The Common Configuration for Core Services in Jazz for Service Management configuration parameters are displayed if IBM Dashboard Application Services Hub is already installed. You must specify an IBM Dashboard Application Services Hub administrative user name and password, such as smadmin and its password. Click **Validate** to confirm the credentials can be used to log in to the dashboard hub. These credentials are used to deploy the Authorization Policy Server application into the WebSphere Application Server for IBM Dashboard Application Services Hub. This user ID is also assigned to the PolicyAdministrator role that is used by the Authorization Policy Server to control which users can create and work with authorization policies.
- The Restart Dashboard Application Services Hub configuration parameters allow you to specify how to restart the Dashboard Application Services Hub after the authorization policy server is installed and configured. Select either an automatic or manual restart of the Dashboard Application Services Hub and click **Confirm**. Click **OK** when prompted to confirm your selection.
- The advanced parameters for the Tivoli Authorization Policy package allow you to specify audit log properties and how to handle distribution of authorization policies. The Tivoli Authorization Policy Server periodically compresses a file of the current set of authorization policies that is

available for distribution. On a periodic interval, the dashboard data provider component of the portal server makes a request to the Tivoli Authorization Policy Server for the latest compressed file of policies. If there is a new file, it is obtained and extracted and this set of policies becomes the current set of policies that are used by the dashboard data provider.

Configure the Authorization Policy Server audit and policy distribution properties. Each property has the following default value and range:

Audit log file count

The maximum number of audit log files to keep at one time.

Default value is 5. Range is greater than 1 and less than 99999.

Audit log file size (megabytes)

The maximum size of each log file in megabytes.

Default size is 10. Range is greater than 1 and less than 99999.

Audit log file directory

The directory into which the log files are stored.

Default value is <JAZZSM_INSTALL_DIR>\AuthPolicyServer\PolicyServer\audit

Policy-distribution Polling interval (minutes)

This property specifies how often the Authorization Policy Server updates the compressed file containing the authorization policies that is downloaded by the dashboard data provider.

Default value is 5. Range is 1 - 1440 minutes.

Policy-distribution Polling directory

The directory into which the version of the policies for distribution is stored.

Default value is <JAZZSM_INSTALL_DIR>\AuthPolicyServer\PolicyServer\dist

- If you are also installing the tivcmd CLI package, configure the IBM GSKit Security parameters. Specify the encryption key that is being used in your IBM Tivoli Monitoring environment in the GSKit field. The default value is *IBMTivoliMonitoringEncryptionKey*.

10. After specifying all of the configuration parameters, click **Next**.

Note: The **Next** button is not enabled until all required configuration parameters are completed.

11. Review the information on the Install Packages Summary window. This is your opportunity to click **Back** and make any changes before proceeding.
12. Click **Install** to begin installing. After installation is complete, the packages are shown in a panel with a message confirming that they are installed. You can now review the log files.
13. Click **Finish** to finalize the installation.

What to do next

See “Verifying the Tivoli Authorization Policy Server installation” on page 329 for steps to verify that the installation was successful.

If you only selected the installation feature of the Tivoli Authorization Policy Server, review the authorization policy server's deployment scripts and then use the **Modify** operation of IBM Installation Manager to deploy the authorization policy server into IBM Dashboard Application Services Hub. See “Configuring the Tivoli Authorization Policy Server feature after installation” on page 328.

If you selected both the installation and configuration features of Authorization Policy Server, perform the following tasks:

1. If you selected the option for a manual restart of Dashboard Application Services Hub in step 9, you must restart the application server for Dashboard Application Services Hub. You must perform this action before you can use the Authorization Policy Server with the tivcmd CLI to create authorization policies.
2. Install the tivcmd Command Line Interface on the computers that administrators use to create authorization policies. Follow the instructions in “Installing the tivcmd Command Line Interface for Authorization Policy using the graphical user interface” on page 330 or “Installing the tivcmd Command Line Interface for Authorization Policy using console mode” on page 332. The user credentials that you specified during the Authorization Policy Server installation are assigned to the PolicyAdministrator role. You must use these credentials with the tivcmd Command Line Interface to log in to the Authorization Policy Server and assign other administrators permission to create and work with authorization policy roles.

The *IBM Tivoli Monitoring Administrator's Guide* provides examples of creating authorization policies for common scenarios in the *Using role-based authorization policies* chapter. Steps for configuring a dashboard environment to use authorization policies are outlined in the *Preparing your dashboard environment* chapter. For a complete list of tivcmd CLI commands, see the *IBM Tivoli Monitoring Command Reference*.

Installing and configuring the Tivoli Authorization Policy Server using console mode

You can use IBM Installation Manager in console mode to install packages from a command-line interface.

Before you begin

Install the prerequisite software that is described in “Required software and memory requirements for a dashboard environment” on page 91 before you install and configure the Tivoli Authorization Policy Server.




Set up a local or network repository that contains the IBM Tivoli Monitoring components for a dashboard environment. See “Using IBM Installation Manager to install or update components from a local repository” on page 300 or “Using IBM Installation Manager to install or update components from a network repository” on page 302.

About this task

Complete the following steps to run the IBM Installation Manager in console mode to install Tivoli Authorization Policy Server on the computer where IBM Dashboard Application Services Hub is already installed.

Note: You cannot install the Authorization Policy Server and the tivcmd Command Line Interface for Authorization Policy at the same time by using console mode. You must install these software packages separately. However, you can install the IBM Infrastructure Management Dashboards for Servers at the same time as the authorization policy server.

Procedure

1. Log in as the user who installed IBM Application Services Hub.
2. Start the IBM Installation Manager in console mode from the directory where it was installed. For example:
 -  `C:\Program Files\IBM\Installation Manager\eclipse\tools\imcl.exe -c`
 -   `/opt/IBM/InstallationManager/eclipse/tools/imcl -c`
3. The following menu is displayed:

```

Select:
  1. Install - Install software packages
  2. Update - Find and install updates and fixes to installed software packages
  3. Modify - Change installed software packages
  4. Roll Back - Revert to an earlier version of installed software packages
  5. Uninstall - Remove installed software packages

Other Options:
  L. View Logs
  S. View Installation History
  V. View Installed Packages
  -----
  P. Preferences
  -----
  E. Export Data for Problem Analysis
  A. About IBM Installation Manager
  -----
  X. Exit Installation Manager

```

Select **1** to install software packages.

4. When you are presented with the list of software packages, select the option for the **Tivoli Authorization Policy Server** package.

Note: You can also select the IBM Infrastructure Management Dashboards for Servers package if you are installing it on the same computer.

5. Select **1** to install the current version of the software package. Option **2** displays all available versions of the package.
6. Authorization Policy Server is selected by default. Select **N** (Next) to continue.
7. Select your preferred installation location and then select **N** (Next) to continue.
8. Select the Tivoli Authorization Policy Server features to install. By default, both *Installation* and *Configuration* are selected. Select **N** (Next) to continue.

Note: It is recommended to select the Tivoli Authorization Policy Server *Installation* and *Configuration* features together (which is the default behavior). However, you cannot select the authorization policy server Configuration feature if you did not select the Configuration feature when IBM Dashboard Application Services Hub was installed. You should also not select the Configuration feature if you want to review the deployment scripts before they are used to deploy the authorization policy server into IBM Dashboard Application Services Hub. If you select the Installation feature on its own, only the authorization policy server binaries and install scripts are laid down in the installation directory (on Windows this location is usually C:\Program Files\IBM\JazzSM\AuthPolicyServer and on Linux/UNIX this location is usually /opt/IBM/JazzSM/AuthPolicyServer). If you select the Configuration feature on its own, the Installation feature is selected as well.

9. Enter your WebSphere administrator-level user credentials (user name and password) that can be used to install the Tivoli Authorization Policy Server application into IBM Dashboard Application Services Hub. The installer validates the entered credentials.
10. The Tivoli Authorization Policy Server periodically compresses a file of the current set of authorization policies that is available for distribution. On a periodic interval, the dashboard data provider component of the portal server makes a request to the Tivoli Authorization Policy Server for the latest compressed file of policies. If there is a new file, it is obtained and extracted and this set of policies becomes the current set of policies that are used by the dashboard data provider.

Enter the Authorization Policy Server advanced configuration values or press **Enter** to select the default values. Each property has the following default value and range:

Audit log file count

The maximum number of audit log files to keep at one time.

Default value is 5. Range is greater than 1 and less than 99999.

Audit log file size (megabytes)

The maximum size of each log file in megabytes.

Default size is 10. Range is greater than 1 and less than 99999.

Audit log file directory

The directory into which the log files are stored.

Default value is <JAZZSM_INSTALL_DIR>\AuthPolicyServer\PolicyServer\audit

Policy-distribution Polling interval (minutes)

This property specifies how often the Authorization Policy Server updates the compressed file that contains the authorization policies that is downloaded by the dashboard data provider.

Default value is 5. Range is 1 - 1440 minutes.

Policy-distribution Polling directory

The directory into which the version of the policies for distribution is stored.

Default value is <JAZZSM_INSTALL_DIR>\AuthPolicyServer\PolicyServer\dist

11. You must restart the Dashboard Application Services Hub after the Authorization Policy Server installation and configuration is complete. Select either an automatic or manual restart of the Dashboard Application Services Hub.
12. Review the installation summary window and select **I** to proceed with the installation.

What to do next

See “Verifying the Tivoli Authorization Policy Server installation” on page 329 for steps to verify that the installation was successful.

If you only selected the installation feature of the Tivoli Authorization Policy Server, review the authorization policy server's deployment scripts and then use the **Modify** operation of IBM Installation Manager to deploy the authorization policy server into IBM Dashboard Application Services Hub. See “Configuring the Tivoli Authorization Policy Server feature after installation” on page 328.

If you selected both the installation and configuration features of Authorization Policy Server, perform the following tasks:

1. If you selected the option for a manual restart of Dashboard Application Services Hub in step 11, you must restart the application server for Dashboard Application Services Hub before the Authorization Policy Server can be used with the tivcmd CLI to create authorization policies.
2. Install the tivcmd Command Line Interface on the computers that administrators use to create authorization policies. Follow the instructions in “Installing the tivcmd Command Line Interface for Authorization Policy using the graphical user interface” on page 330 or “Installing the tivcmd Command Line Interface for Authorization Policy using console mode” on page 332. The user credentials that you specified during the Authorization Policy Server installation are assigned to the PolicyAdministrator role. You must use these credentials with the tivcmd Command Line Interface to log in to the Authorization Policy Server and assign other administrators permission to create and work with authorization policy roles.

The *IBM Tivoli Monitoring Administrator's Guide* provides examples of creating authorization policies for common scenarios in the *Using role-based authorization policies* chapter. Steps for configuring a dashboard environment to use authorization policies are outlined in the *Preparing your dashboard environment* chapter. For a complete list of tivcmd CLI commands, see the *IBM Tivoli Monitoring Command Reference*.

Configuring the Tivoli Authorization Policy Server feature after installation

If you only chose the Installation feature for the Tivoli Authorization Policy Server package at installation time, use the Modify operation of the IBM Installation Manager's to select the Configuration feature of the package after you have reviewed the Authorization Policy Server's deployment scripts. When you select the authorization policy server package, it is deployed to IBM Dashboard Application Services Hub.

Before you begin

Ensure that the configuration feature for IBM Dashboard Application Services Hub is already selected. The Tivoli Authorization Policy Server is an extension offering of the Dashboard Application Services Hub. As a result, an error message is displayed if you attempt to select the Authorization Policy Server configuration feature when Dashboard Application Services Hub is not configured.

About this task

If you selected the Installation feature when installing the packages, only the binaries and install scripts are laid down in your installation directory:

- **Windows** C:\Program Files\IBM\JazzSM\AuthPolicyServer
- **Linux** **UNIX** /opt/IBM/JazzSM/AuthPolicyServer

To add the Configuration feature, you must perform the following steps to modify your Tivoli Authorization Policy Server package:

Procedure

1. Log in as the user who installed IBM Dashboard Application Services Hub and authorization policy server.
2. Start the IBM Installation Manager:
 - **Windows** Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager.
 - **Linux** **UNIX** Execute the IBMIM binary under /opt/IBM/InstallationManager/eclipse.
3. On the main Installation Manager window, click **Modify**.
4. Select **Core services in Jazz for Service Management**. This contains the Tivoli Authorization Policy Server package. Click **Next** to continue.
5. Expand the **Authorization Policy Server** node and select the **Configuration** check box. Click **Next** to continue.
6. Specify an IBM Dashboard Application Services Hub administrative user name and password, such as smadmin and its password for the **Common Configurations for Core services in Jazz for Service Management** configuration parameters, and click **Validate**. These credentials are used to deploy the **Infrastructure Management Dashboards for Servers** application into the WebSphere Application Server for the IBM Dashboard Application Services Hub.
7. The Tivoli Authorization Policy Server periodically compresses a file of the current set of authorization policies that is available for distribution. On a periodic interval, the dashboard data provider component of the portal server makes a request to the Tivoli Authorization Policy Server for the latest compressed file of policies. If there is a new file, it is obtained and extracted and this set of policies becomes the current set of policies that are used by the dashboard data provider.
Configure the Authorization Policy Server audit and policy distribution properties. Each property has the following default value and range:

Audit log file count

The maximum number of audit log files to keep at one time.

Default value is 5. Range is greater than 1 and less than 99999.

Audit log file size (megabytes)

The maximum size of each log file in megabytes.

Default size is 10. Range is greater than 1 and less than 99999.

Audit log file directory

The directory into which the log files are stored.

Default value is <JAZZSM_INSTALL_DIR>\AuthPolicyServer\PolicyServer\audit

Policy-distribution Polling interval (minutes)

This property specifies how often the Authorization Policy Server updates the compressed file that contains the authorization policies that is downloaded by the dashboard data provider.

Default value is 5. Range is 1 - 1440 minutes.

Policy-distribution Polling directory

The directory into which the version of the policies for distribution is stored.

Default value is <JAZZSM_INSTALL_DIR>\AuthPolicyServer\PolicyServer\dist

8. You must restart the Dashboard Application Services Hub after the Tivoli Authorization Policy Server installation is complete. Select either an automatic or manual restart of the Dashboard Application Services Hub and click **Confirm**. Click **OK** when prompted and **Next** to continue.
9. The Modify Packages summary panel is displayed. Verify the changes and click **Modify** to add the Configuration feature.
10. The Results window displays the package modification status. Click **Finish** to complete the modification.

What to do next

If you selected the option for a manual restart of Dashboard Application Services Hub in step 8, you must restart the application server for Dashboard Application Services Hub. You must restart the application server before you can use the Authorization Policy Server with the tivcmd CLI to create authorization policies.

Install the tivcmd Command Line Interface on the computers that administrators use to create authorization policies. Follow the instructions in “Installing the tivcmd Command Line Interface for Authorization Policy using the graphical user interface” on page 330 or “Installing the tivcmd Command Line Interface for Authorization Policy using console mode” on page 332. The user credentials that you specified during the Authorization Policy Server installation are assigned to the PolicyAdministrator role. You must use these credentials with the tivcmd Command Line Interface to log in to the Authorization Policy Server and assign other administrators permission to create and work with authorization policy roles.

The *IBM Tivoli Monitoring Administrator's Guide* provides examples of creating authorization policies for common scenarios in the *Using role-based authorization policies* chapter. Steps for configuring a dashboard environment to use authorization policies are outlined in the *Preparing your dashboard environment* chapter. For a complete list of tivcmd CLI commands, see the *IBM Tivoli Monitoring Command Reference*.

Verifying the Tivoli Authorization Policy Server installation

If you selected both the *Installation* and *Configuration* features when installing the packages, the binaries and installation scripts are laid down and the Authorization Policy Server is deployed in the enterprise archive file of the Dashboard Application Services Hub.

Perform the following checks to verify the Tivoli Authorization Policy Server installation:

- Check that there are no errors in the deployPolicyServer.log file that is located in `<install_dir>\JazzSM\logs`.
- Check that the AuthPolicyServer folder is in `<install_dir>\JazzSM` and that the installation binaries are present.
- Check that the PolicyServer.war folder is in `<install_dir>\JazzSM\profile\installedApps\localhostNode01Cell\isc.ear\`.

If you selected the *Installation* feature only when installing the packages, the binaries and installation scripts are laid down and the Authorization Policy Server is not deployed. Perform the following check to confirm this:

- Check that the AuthPolicyServer folder is in `<install_dir>\JazzSM` and that the installation binaries are present.

Installing the tivcmd Command Line Interface for Authorization Policy using the graphical user interface

Use the IBM Tivoli Monitoring Installation Launchpad or IBM Installation Manager with a local or network repository to install the tivcmd Command Line Interface for Authorization Policy using a graphical user interface. The Installation Launchpad is the recommended approach if you do not have experience using IBM Installation Manager.

Before you begin

The tivcmd CLI should be installed on the computers that administrators use to create and work with authorization policies. The tivcmd CLI connects to an IBM Dashboard Application Services Hub where the Tivoli Authorization Policy Server package is installed and configured using HTTP or HTTPS. Therefore, it can be installed on a computer without the Dashboard Application Services Hub and authorization policy server. You can install the tivcmd CLI on the same computer as the authorization policy server.

If you are installing the tivcmd CLI on a computer without the authorization policy server and Dashboard Application Services Hub, follow the procedure below.

If you want to install the tivcmd CLI on the same computer as the authorization policy server using a graphical user interface, you can install both components at the same time by using the instructions in “Installing and configuring the Tivoli Authorization Policy Server using the graphical user interface” on page 320.

If you are using the IBM Tivoli Monitoring Installation Launchpad to perform the installation, there are no prerequisites to install.

If you are using an IBM Installation Manager repository to perform the installation, you must complete the following setup steps:

1. Install IBM Installation Manager from the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image, if it is not already installed or is at a version before V1.6.1. IBM Installation Manager version 1.6.1 is provided on the DVD under the `im.<operating system>.<architecture>` directories. For more details about performing the installation, see the IBM Installation Manager information center at <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.
2. Create a repository containing the IBM Tivoli Monitoring components for a dashboard environment. See “Using IBM Installation Manager to install or update components from a local repository” on page 300 or “Using IBM Installation Manager to install or update components from a network repository” on page 302.

About this task

Use the following steps to install the tivcmd Command Line Interface for Authorization Policy.

Procedure

1. If you are using the IBM Tivoli Monitoring Installation Launchpad, perform these steps to start the installation:
 - a. Log in to the system as the user who will run the tivcmd CLI.
 - b. Execute the launchpad command or script for your operating system. The launchpad commands and scripts are located in the root directory of the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components DVD* or DVD image.
 - **Windows** 32 bit system launchpad.exe
 - **Windows** 64 bit system launchpad64.exe
 - **Linux** **UNIX** launchpad.sh
 - c. On the main launchpad panel, select one of the following links under the **Product Overview** category:
 - *Install as administrative user*: Select this link if you want to install and run the tivcmd CLI as an administrative user on Windows or as root on Linux/UNIX.
 - *Install as non-administrative user*: Select this link if you do not want to install and run the tivcmd CLI as an administrative user on Windows or as root on Linux/UNIX.
 - d. In the Installation Manager main window, click **Install** to display the available packages.
2. If you are using IBM Installation Manager, perform the following steps to start the installation:
 - a. Log in to the computer as the user who will run the tivcmd CLI commands.
 - b. Start the IBM Installation Manager:
 - **Windows** Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager.
 - **Linux** **UNIX** Execute the IBMIM binary under /opt/IBM/InstallationManager/eclipse.

The IBM Installation Manager window opens with several choices for managing your IBM software installations.
 - c. In the Installation Manager main window, click **Install** to see the available packages in the repositories that are configured for IBM Installation Manager.
3. Select the **tivcmd Command Line Interface for Authorization Policy** package and click **Next**.

Note: If you are using the IBM Tivoli Monitoring Installation Launchpad, the **tivcmd Command Line Interface for Authorization Policy** package is preselected.

4. If you are installing other packages at the same time, you might be asked to accept a license agreement for the other packages. After you accept the agreement, click **Next**.
5. If you are using IBM Installation Manager to perform the installation, specify the location of the Shared Resources Directory and click **Next**.

You are asked to specify a shared resources directory the first time that you install a package on a computer. The shared resources directory is where IBM Installation Manager stores installation artifacts that can be used by one or more package groups. Choose a directory on your largest drive. You cannot change the shared resources directory location until you uninstall all packages.
6. Configure the installation directory for the tivcmd Command Line Interface package group. The default installation directory is as follows:
 - **Windows** C:\IBM\TivoliMonitoring
 - **Linux** **UNIX** /opt/IBM/TivoliMonitoring.

Note:

- a. You cannot install the tivcmd CLI under the same installation directory as Jazz for Service Management or the monitoring server, portal server, portal client, monitoring agents, and tacmd CLI components of IBM Tivoli Monitoring.
- b. On a 64 bit machine, the 64 bit architecture of the tivcmd Command Line Interface for Authorization Policy is installed by default. You can select to switch to the 32 bit architecture. However, on zSeries systems with Linux, you can only install the 32 bit architecture of the CLI.

Click **Next** to continue.

7. If you are installing other packages at the same time, you might be asked to select the translation packages to install for those packages. After you make your selection, click **Next**.

Note: The language files for all supported languages are always installed with the tivcmd Command Line Interface for Authorization Policy package.

8. Make sure that the **tivcmd Command Line Interface for Authorization Policy** feature is selected and click **Next**.
9. Configure the IBM GSKit Security parameters. Specify the encryption key that is used in your IBM Tivoli Monitoring environment in the GSKit field. The default value is *IBMTivoliMonitoringEncryptionKey*. Click **Next** to continue.
10. Review the information on the Install Packages Summary window. This is your opportunity to click **Back** and make any changes before proceeding.
11. Click **Install** to begin installing. After installation is complete, the packages are shown in a panel with a message confirming that they are installed. At this point, you can review the log files.
12. The Results window displays the package installation status. Click **Finish** to complete the installation.

What to do next

Use the tivcmd Command Line Interface to create authorization policies. The *IBM Tivoli Monitoring Administrator's Guide* provides examples of creating authorization policies for common scenarios in the *Using role-based authorization policies* chapter. Steps for configuring a dashboard environment to use authorization policies are outlined in the *Preparing your dashboard environment* chapter. For a complete list of tivcmd CLI commands, see the *IBM Tivoli Monitoring Command Reference*.

Installing the tivcmd Command Line Interface for Authorization Policy using console mode

You can use IBM Installation Manager in console mode to install packages from a command line interface.

Before you begin

Install IBM Installation Manager from the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image, if it is not already installed or is at a version before V1.6.1. IBM Installation Manager version 1.6.1 is provided on the DVD under the *im.<operating system>.<architecture>* directories. For more details about performing the installation, see the IBM Installation Manager information center at <http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp>.

Create a repository containing the IBM Tivoli Monitoring components for a dashboard environment. See “Using IBM Installation Manager to install or update components from a local repository” on page 300 or “Using IBM Installation Manager to install or update components from a network repository” on page 302.

About this task

Perform the following steps to run the IBM Installation Manager in console mode to install the tivcmd CLI.

Note: You cannot install the Authorization Policy Server and the tivcmd Command Line Interface for Authorization Policy at the same time using console mode. These software packages must be installed separately.

Procedure

1. Log in as the user who installed Tivoli Authorization Policy Server if it is already installed on this computer. Otherwise, log in as the user who will run the tivcmd CLI commands.
2. Start the IBM Installation Manager in console mode from the directory where it was installed. For example:
 - **Windows** `C:\Program Files\IBM\Installation Manager\eclipse\tools\imcl.exe -c`
 - **Linux** **UNIX** `/opt/IBM/InstallationManager/eclipse/tools/imcl -c`
3. The following menu is displayed:

```
Select:
  1. Install - Install software packages
  2. Update - Find and install updates and fixes to installed software packages
  3. Modify - Change installed software packages
  4. Roll Back - Revert to an earlier version of installed software packages
  5. Uninstall - Remove installed software packages

Other Options:
  L. View Logs
  S. View Installation History
  V. View Installed Packages
  -----
  P. Preferences
  -----
  E. Export Data for Problem Analysis
  A. About IBM Installation Manager
  -----
  X. Exit Installation Manager
```

Select **1** to install software packages.

4. When you are presented with the list of software packages, select the option for the **tivcmd Command Line Interface for Authorization Policy** package.
5. Select **1** to install the current version of the software package. Option **2** displays all available versions of the package.
6. The tivcmd Command Line Interface for Authorization Policy package is selected by default. Select **N** (Next) to continue.

You are asked to specify a shared resources directory the first time that you install a package on a computer system. The shared resources directory is where IBM Installation Manager stores installation artifacts that can be used by one or more package groups. Choose a directory on your largest drive. You cannot change the shared resources directory location until you uninstall all packages.
7. Next, you are presented with options for changing the location of the tivcmd CLI installation directory and the tivcmd CLI architecture. Select **N** (Next) to accept the defaults or choose an option to change the installation directory or architecture:
 - a. Select **M** to change the installation directory for the tivcmd Command Line Interface package group. The default installation directory is **Windows** `C:\IBM\Tivoli Monitoring` **Linux** **UNIX** `/opt/IBM/TivoliMonitoring`.

Note: The tivcmd CLI cannot be installed under the same installation directory as Jazz for Service Management or the monitoring server, portal server, portal client, monitoring agents, and tacmd CLI components of IBM Tivoli Monitoring.

- b. Select **T** to change the architecture type.

On a 64 bit machine, the 64 bit architecture of the tivcmd Command Line Interface for Authorization Policy is installed by default. You can select to switch to the 32 bit architecture. On zSeries systems with Linux, you can only install the 32 bit architecture of the tivcmd CLI.

- c. Select **N** (Next) to continue.

8. The **tivcmd Command Line Interface for Authorization Policy** feature is selected by default. Select **N** (Next) to continue.
9. Either use the default GSKit key value of *IBMTivoliMonitoringEncryptionKey* or enter the GSKit key value used in your IBM Tivoli Monitoring environment and then press **Enter** to continue.
10. Select **N** (Next) and press Enter to continue.
11. Review the installation summary window and select **I** to proceed with the installation.

What to do next

Use the tivcmd Command Line Interface to create authorization policies. The *IBM Tivoli Monitoring Administrator's Guide* provides examples of creating authorization policies for common scenarios in the *Using role-based authorization policies* chapter. Steps for configuring a dashboard environment to use authorization policies are outlined in the *Preparing your dashboard environment* chapter. For a complete list of tivcmd CLI commands, see the *IBM Tivoli Monitoring Command Reference*.

Updating the Tivoli Authorization Policy Server using the graphical user interface

Use the IBM Installation Manager graphical user interface to update the Tivoli Authorization Policy Server to a later version such as a Fix Pack version.

Before you begin




See “Required software and memory requirements for a dashboard environment” on page 91 to determine if you need to update any of the prerequisite software before updating Tivoli Authorization Policy Server.

Set up a local or network repository that contains the IBM Tivoli Monitoring components for a dashboard environment. See “Using IBM Installation Manager to install or update components from a local repository” on page 300 and “Using IBM Installation Manager to install or update components from a network repository” on page 302.

About this task

Complete the following steps to update the Tivoli Authorization Policy Server using a graphical user interface.

Procedure

1. Start the IBM Installation Manager:
 -  Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager\eclipse.
 -   Execute the IBMIM binary under /opt/IBM/InstallationManager/eclipse.
2. On the main Installation Manager window, click **Update**.
3. Select the **Core services in Jazz for Service Management** package and click **Next**.

4. Installation Manager displays the version of the updated Tivoli Authorization Policy Server package. Click **Next** to continue.
5. Select the Installation and Configuration features of the Tivoli Authorization Policy Server package and click **Next**.

Note: It is recommended to select the **Tivoli Authorization Policy Server Installation** and **Configuration** features together (which is the default behavior). You should not select the **Configuration** feature if you want to review the deployment scripts before they are used to update the authorization policy server in IBM Dashboard Application Services Hub. If you select the **Installation** feature on its own, only the updated authorization policy server binaries and install scripts are laid down in the installation directory (on Windows this location is usually C:\Program Files\IBM\JazzSM\AuthPolicyServer and on Linux/UNIX this location is usually /opt/IBM/JazzSM/AuthPolicyServer) and you are not prompted for the configuration parameters in the steps that follow. If you select the **Configuration** feature on its own, the **Installation** is also selected.

6. The WebSphere configuration panel is displayed. Enter your administrator-level user credentials, such as smadmin and its password in the **User name** and **Password** fields, and click **Validate**. These credentials are used to update the authorization policy server application in the WebSphere Application Server for IBM Dashboard Application Services Hub. After your credentials are validated, the **Next** button is enabled. Click **Next**.
7. The Tivoli Authorization Policy Server periodically compresses a file of the current set of authorization policies that is available for distribution. On a periodic interval, the dashboard data provider component of the portal server makes a request to the Tivoli Authorization Policy Server for the latest compressed file of policies. If there is a new file, it is obtained and extracted and this set of policies becomes the current set of policies that are used by the dashboard data provider.

Configure the advanced Tivoli Authorization Policy Server parameters for audit logging and policy distribution. Each property has the following default value and range:

Audit log file count

The maximum number of audit log files to keep at one time.

Default value is 5. Range is greater than 1 and less than 99999.

Audit log file size (megabytes)

The maximum size of each log file in megabytes.

Default size is 10. Range is greater than 1 and less than 99999.

Audit log file directory

The directory into which the log files are stored.

Default value is <JAZZSM_INSTALL_DIR>\AuthPolicyServer\PolicyServer\audit

Policy-distribution Polling interval (minutes)

This property specifies how often the Authorization Policy Server updates the compressed file containing the authorization policies that is downloaded by the dashboard data provider.

Default value is 5. Range is 1 - 1440 minutes.

Policy-distribution Polling directory

The directory into which the latest version of the policies for distribution is stored.

Default value is <JAZZSM_INSTALL_DIR>\AuthPolicyServer\PolicyServer\dist

Note: The panel displays the values that were specified the last time that the Tivoli Authorization Policy package was configured for an installation or package update using the IBM Tivoli Monitoring installation launchpad or IBM Installation Manager. If you modified the properties after the last configuration, then those values are not displayed. If you want to continue using the modified values, you must reenter them on this panel.

8. You must restart the Dashboard Application Services Hub after the Tivoli Authorization Policy Server update is complete. Select either an automatic or manual restart of the Dashboard Application Services Hub and click **Confirm**. Click **OK** when prompted and **Next** to continue.
9. Click **Update** to update the Tivoli Authorization Policy Server.

What to do next

If you only selected the Installation feature of the updated Tivoli Authorization Policy Server package, review the authorization policy server's deployment scripts and then use the **Modify** operation of IBM Installation Manager to deploy the updated authorization policy server into IBM Dashboard Application Services Hub. For more information, see “Configuring the Tivoli Authorization Policy Server feature after installation” on page 328.

If you selected both the installation and configuration features of Authorization Policy Server together with the option for a manual restart of Dashboard Application Services Hub in step 9, you must restart the application server for Dashboard Application Services Hub. You must do this before you can use the updated Authorization Policy Server with the `tivcmd` CLI to create authorization policies.

Rollback of the Tivoli Authorization Policy Server

You can use the rollback feature of IBM Installation Manager to revert an updated package back to a previous version. To complete the rollback, the Installation Manager requires information from the earlier version of the package, if the package supports rollback.

Before you begin

When Installation Manager rolls a package back to a previous version, it uninstalls the current version of some files and reinstalls the previous version of those files.

By default, the files that are required to complete the rollback are saved on your computer. On the **Files for Rollback** preference page, you can change the default action and remove any previously stored files. For more information, see “Saving package information for rollback” on page 305.

About this task

Complete the following steps to rollback the Tivoli Authorization Policy Server if IBM Installation Manager was configured to save the files for the previous version of the package:

Procedure

1. On the main Installation Manager window, click **Rollback** to start the Rollback wizard.
2. From the **Package Group Name** list, select the package group that contains the Authorization Policy Server package that you want to rollback and click **Next**.
3. Select the version of the package that you want to rollback to and click **Next**.
4. Review the summary information and click **Rollback** to rollback the package.
5. Follow the steps that are displayed in the wizard. These are the same steps as “Updating the Tivoli Authorization Policy Server using the graphical user interface” on page 334.




Installing domain definitions for Tivoli Performance Analyzer

Install the domain definitions to enable performance analytics for a broader set of systems including DB2, Oracle, System p Series, ITCAM for RT, and VMware. You can install domain definitions in GUI mode or console mode. The domain definitions for the Tivoli Performance Analyzer are found on the IBM Tivoli Performance Analyzer Reports and Domain Definitions DVD or DVD image.

Note:

1. Tivoli Performance Analyzer is a special agent that works on data collected by other agents. Tivoli Performance Analyzer is shipped with separate packages called domains. These domains bundle together a set of artifacts that are used to process and present warehoused data collected by other agents. For example, there is Performance Analyzer DB2 domain, Performance Analyzer OS Domain, Performance Analyzer Oracle Domain, and so on.
2. Domain definitions for Tivoli Performance Analyzer must be added in order to complete the product configuration. You must install the domains on the server that has the hub Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, and Tivoli Performance Analyzer installed. In a distributed environment, you must perform domain installation on each server where at least one of these components is installed: hub Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, and Tivoli Performance Analyzer.
3. When run on a host, the domain installer presents options to add domain support for only those components that it finds installed on the server locally, for example Tivoli Enterprise Portal Server, Tivoli Enterprise Portal clients - Browser and Desktop, Tivoli Enterprise Monitoring Server, or Tivoli Performance Analyzer itself. Choose the installed component as appropriate.
4. Due to new domain product codes introduced with Tivoli Monitoring V6.2.3, any historical data collection settings for the Tivoli Performance Analyzer are not maintained when upgrading from V6.2.2 to V6.2.3 or higher. The domain installer image includes scripts that create historical data collection settings. You can use these scripts to quickly recreate any historical data collection settings that were lost.

The logs for the domain installation can be found at:

-  `<ITM DIR>\logs\itpa_domain.log`
-   `<ITM DIR>/logs/itpa_domain.log`

Installing domain definitions in GUI mode

Use the GUI to install the domain definitions automatically and check if all the required programs for configuration are in place.

About this task

To install the required domain definitions:

Procedure

1. From the product CD launch `setup.exe`. The installer window opens.
2. Click **Next** on the installer welcome page.
3. In the **License Agreement** window, read the agreement, accept the terms and click **Next**.
4. Select the components for which you want to add domain definitions and click **Next**.

Note: If you do not mark all the components, there may be a lack of coherence between the selected ones.

If you select **Tivoli Performance Analyzer Domain support**, only the **active** or **inactive** state of tasks for the selected domain is preserved after reinstallation in the upgrade scenario. All other attributes are set back to **Default**.

Domains require Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Tivoli Enterprise Portal Desktop, Tivoli Enterprise Portal Browser and Tivoli Performance Analyzer support files to be installed.

5. Review the settings and click **Next**.
6. Select the configuration type and click **Next**.
7. Enter the host name of the machine where the Tivoli Enterprise Portal Server resides and click **Next**.
8. Choose where to locate the application support for the Tivoli Enterprise Monitoring Server and click **OK**.

9. In the next window, select the application support you want to add to the Tivoli Enterprise Monitoring Server. Choose the component and add or update the situations and click **OK**.
10. In order to finish the installation, check the box to see the installation documentation and click **Finish**.

Results

You have now automatically broadened the Tivoli Performance Analyzer domain support.

What to do next

Important: Restart Tivoli Monitoring components after the installation. The tool shows which components need to be restarted.

Installing the domain definitions in console mode

Use the console mode to install the domain definitions in an interactive manner.

Before you begin

Tivoli Performance Analyzer Agent and its supports must be installed before installing any domain definitions.

About this task

Use the following steps to install the required domain definitions on a UNIX computer:

Procedure

1. In the directory where you extracted the installation files, run the following command:
`./install.sh`
2. When prompted for the IBM Tivoli Monitoring home directory, press **Enter** to accept the default (/opt/IBM/ITM). If you want to use a different installation directory, type the full path to that directory and press **Enter**.
3. If the directory you specified does not exist, you are asked whether to create it. Type 1 to create this directory.
4. Type 1 in order to choose **Install products to the local host** from the prompt and start installation.
5. The user license agreement is displayed. Press **Enter** to read through the agreement.
6. Type 1 to accept the agreement and press **Enter**.
7. Select a locally installed component for which you want to add domain definitions.

Note:

- a. If you select **Tivoli Performance Analyzer Domain support**, only the **active** or **inactive** state of tasks for the selected domain is preserved after reinstallation in the upgrade scenario. All other attributes are set back to **Default**.
 - b. Domains require Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Tivoli Enterprise Portal Desktop, Tivoli Enterprise Portal Browser and Tivoli Performance Analyzer support files to be installed.
8. To select more components for which you want to add domain definitions, type **1**, which takes you to step 5. To finish the installation, select **2**.

Results

You have now installed the required domain definitions and exited the console mode.

Attention: In order to display workspaces properly, reconfigure Tivoli Enterprise Portal Server and Tivoli Enterprise Portal Desktop Client. You also must restart Tivoli Performance Analyzer to start calculations for newly installed analytical tasks.

Installing language packs

Language support for products for which application support is provided with IBM Tivoli Monitoring appears on the following media.

- *IBM Tivoli Monitoring Language Support* : Brazilian Portuguese, French, German, Italian, Japanese, Korean, Simplified Chinese, Spanish, Traditional Chinese, Czech, Hungarian, Polish, Russian, Thai
- *IBM Tivoli Monitoring Agent Builder Toolkit Language Support DVD*
- *IBM Tivoli Monitoring Agents for System P Language Support CD*
- *IBM Tivoli Monitoring Log File Monitoring Agent DVD*
- Agent product installation CDs

The IBM Tivoli Monitoring Language Support DVDs contain the national language versions of the help and presentation files for the components and agents shipped with IBM Tivoli Monitoring.




The IBM Infrastructure Management Dashboards for Servers, Tivoli Authorization Policy Server, and Tivoli Authorization Policy Command-Line Interface include language support with the base installation image. There is no separate Language pack for those components.

Note: The Tivoli Monitoring self-describing agent feature integrates the installation of an agent with the dispersal and installation of associated product support files throughout your IBM Tivoli Monitoring infrastructure. Language Pack installation is not supported for products installed using self-describing agent capability. For these products you must install language packs manually, using the steps described in this section.

Install the language packs on any system where you have installed the Tivoli Enterprise Portal or where you have installed a desktop client. (If you download and run a desktop client using Web Start, you do not need to install the language packs on the local system. They are downloaded from the portal server.) *Before you can install a language pack, you must install the component in English.*

Before installing a language pack, first install the component in English. Also ensure that Java Runtime Environment version 1.5 or above is installed and set in the system path. Perform the following steps to install a language pack on any system where you have installed either the Tivoli Enterprise Portal Server or the Tivoli Enterprise Portal desktop client:

1. In the directory where you extracted the language pack installation image, launch the installation program as follows:

-  Double-click the `lpinstaller.bat` file.
-   Run the following command:
`./lpinstaller.sh -c install_dir`

where:

install_dir

is the directory where you installed IBM Tivoli Monitoring (usually `/opt/IBM/ITM`).

To perform a console installation on Linux or UNIX (instead of a GUI installation), add the `-i` console parameter to the above command.

2. Select the language you want installed, and click OK.
3. On the Introduction panel, click Next.
4. On the Select Action Set panel, click Add/Update, and click Next.
5. Select the folder in which the Language Support package files (win*.jar and unix*.jar) are located, and click Next. The default folder is the directory where the installer is launched.
6. Select the languages that you want to install, and click Next.
For multiple selections, hold down the Ctrl key.
7. Review the installation summary, and, if correct, click Next.
The installation's progress is displayed.
8. On the Post Install Message panel, click Next.
9. Click Done once the installation is complete.
10. Reconfigure and restart the Tivoli Enterprise Portal Server and the Eclipse Help Server. See below.

After installing a Tivoli Monitoring Language Pack, reconfigure the portal server and the desktop client using either the Manage Tivoli Enterprise Monitoring Services utility or the `itmcmd config` command.

Use one of the following methods to reconfigure the affected components:

- Launch Manage Tivoli Enterprise Monitoring Services, right-click the affected component, and select **Reconfigure**. (See “Starting Manage Tivoli Enterprise Monitoring Services” on page 411.)
- Change to the `install_dir/bin` directory, and enter the following commands:

```
./itmcmd config -A cq
./itmcmd config -A cj
```

Accept the default values, which reflect the decisions made when the component was installed or last configured. For instructions on specifying your users' language environment, see the *IBM Tivoli Monitoring Administrator's Guide*.

After you have reconfigured these components, you need to stop and restart these components:

- Tivoli Enterprise Portal Server
- Tivoli Enterprise Portal desktop or browser client

For SuSE Linux Enterprise Server (SLES) 10 computers only: On the SLES 10 platform, the Tivoli Enterprise Portal displays corrupted text resources in the Japanese locale. Download the Kochi fonts contained in the `kochi-substitute-20030809.tar` package from the following website: <http://sourceforge.jp/projects/efont/files/>.

The downloaded tar file includes the truetype fonts (ttf files), which need to be installed on your system. Complete the following steps to install the files:

1. Extract the tar file.
2. Copy the font files (ttf) to X11 font path (for example, `/usr/X11R6/lib/X11/fonts/truetype`).
3. Run **SuSEconfig -module fonts**.

See the following Web site for detailed instructions for installing the additional fonts to SuSE Linux: <http://www.suse.de/~mfabian/suse-cjk/installing-fonts.html>.

Silent installation of language packs

You can install the language pack using a silent installation method.

Procedure

1. Copy and paste the following response file template as `ITM_LP_slient.txt` file.
2. Change the following parameter setting: `NLS_PACKAGE_FOLDER`; `PACKAGE_NAME_LIST`; `LANG_SELECTION_LIST`.

3. Run the following command to silently install the language pack:

- **Windows**
`lpinstaller.exe -f <path_to_response_file>`
- **Linux** **UNIX**
`lpinstaller.sh -c <andle_home> -f <path_to_response_file>`

Where <andle_home> is the IBM Tivoli Monitoring base directory.

Response file template for language pack silent installation:

This section contains the response file template for language pack silent installation.

```
#      IBM Tivoli Monitoring Language Pack Silent Installation Operation
#
#This is a sample response file for silent installation of the IBM Tivoli Monitoring
#Language Pack Installer. The example used in this file is the IBM Tivoli Monitoring
#language pack. This file was built by the Replay feature of InstallAnywhere.
#The variables in this file were set by Panels, Consoles, or Custom Code.
#-----

#-----
#To successfully complete a silent installation of the IBM Tivoli Monitoring
#Language Pack,complete the following steps:
#1.Copy the ITM_LP_silent.txt to the directory where lpinstaller.bat
#or lpinstaller.sh is located
#(IBM Tivoli Monitoring  Language Pack build location).
#
#2.Modify the response file so that it is customized correctly and completely
#for your site.
#Complete all of the following steps in the response file.
#
#3.After customizing the response file, invoke the silent installation using the
#following command:
#For Windows:
#   lpinstaller.exe -f <path_to_response_file>
#For UNIX and Linux:
#   lpinstaller.sh -c <andle_home> -f <path_to_response_file>
#Note:<andle_home> is the ITM base directory

#-----
#Force silent install mode.
#-----
INSTALLER_UI=silent

#-----
#Run add/update actions.
#-----
CHOSEN_INSTALL_SET=ADDUPD_SET

#-----
#Choose a folder.
#Modify the following path to the language package driver location:
#Windows path:
# NLS_PACKAGE_FOLDER=C:\\build_machine\\LP\\ITM622_0723\\tvtbuild\\200909011756
#UNIX path:
NLS_PACKAGE_FOLDER=//windows

#-----
#List the packages to process; both variables are required.
#Each variable requires that full paths are specified.
#Separate multiple entries with a semicolon (;).
#For Windows:
#       Use the backslash-backslash(\\) as a file separator.
#For Unix and Linux:
#       Use the slash-slash (//) as a file separator.
#-----
#Windows path:
```

```
#PACKAGE_NAME_LIST=C:\\build_machine\\LP\\ITM622_0723\\tvtbuild\\200909011756\\winDBCS.jar
#UNIX path:
PACKAGE_NAME_LIST=//windows//unixSBCS.jar;//windows//unixDBCS.jar;//windows//unixCER.jar




#-----
#List the languages to process.
#Separate multiple entries with semicolons.
#-----
LANG_SELECTION_LIST=de;es;it;fr;ja;ko;pt_BR;zh_CN;zh_TW;cs;hu;pl;ru;th
```

Uninstalling a language pack

Take the following steps to uninstall a language pack.

Procedure

1. In the directory where you extracted the language pack installation image, launch the installation program as follows:

-  Double-click the `lpinstaller.bat` file.
-   Run the following command:
`./lpinstaller.sh -c install_dir`

where:

`install_dir`

is the directory where you installed IBM Tivoli Monitoring (usually `/opt/IBM/ITM`).

To perform a console installation on Linux or UNIX (instead of a GUI installation), add the `-i` console parameter to the above command.

2. Select the language to be uninstalled, and click **OK**.
3. On the Introduction panel, click **Next**.
4. On the **Select Action Set** panel, click **Remove**, and click **Next**.
5. Select the languages that you want to uninstall, and click **Next**.
For multiple selections, hold down the Ctrl key.
6. Review the installation summary, and, if correct, click **Next**.
7. On the Post Install Message panel, click **Next**.
8. Click **Done** once the uninstallation is complete.
9. Reconfigure and restart the Tivoli Enterprise Portal Server and the Eclipse Help Server.

Configuring clients, browsers, and JREs

The configuration required for Tivoli Enterprise Portal clients depends upon the client deployment mode being used, the browser being used, the Java runtime environment (JRE) being used, and the operating system the client is being used on.

The following sections discuss the configuration for clients by mode of deployment:

- “Desktop clients” on page 343
- “Browser clients” on page 343
- “Java Web Start clients” on page 347

The version of IBM JRE required by Tivoli Enterprise Portal clients is installed with Tivoli Management Services components. If you want to run a client on a machine where no other components are installed, you can download the IBM JRE installer from the Tivoli Enterprise Portal Server. On Windows, the IBM JRE must be installed as the system JVM if you plan on using the browser client.

The packaging, installation, and servicing of the Oracle JRE is not provided by IBM. The Oracle JRE at version 1.7.0.xx must already be installed on the machines where the Tivoli Enterprise Portal client will run. The Oracle JRE can be downloaded from the following Web site: <http://www.oracle.com/technetwork/java/archive-139210.html>. For help installing the Oracle JRE, visit the following Web site: http://www.java.com/en/download/help/index_installing.xml.

The IBM Tivoli Monitoring installer does not use the system JRE. Instead it installs a local copy of Java, one that is private to Tivoli Monitoring. This applies to all Java-dependent Tivoli Monitoring components.

This embedded JRE is installed in the %CANDLE_HOME%\java directory under Windows, and the %CANDLE_HOME/JRE directory under Linux. The system Java is untouched by the IBM Tivoli Monitoring installer.

Under Windows, the browser client and Java Web Start still use the system JRE. For your convenience, the IBM JRE installation image is distributed as part of the browser client package.

The desktop client uses the embedded JVM by default.

The IBM Tivoli Monitoring installer does not install an embedded JVM if none of the components selected for installation has a Java dependency.

Desktop clients

By default, you do not need to perform any post-installation configuration of the desktop client; however, configuration is required to the start-up script of the desktop client under the following scenarios.

- Both IBM and Oracle JREs are installed on the computer and you want to use the Oracle JRE, or
- You have multiple versions of the Oracle JRE or IBM JRE installed and you want to specify a particular version

On Windows computers, add a user-level environment variable named TEP_JAVA_HOME whose value is the fully qualified directory location of the Oracle JRE you want to use. For example (assuming a Windows 7 desktop), TEP_JAVA_HOME=C:\Program Files (x86)\Java\bin\. On UNIX and Linux computers, define the same environment variable with the fully qualified location of the Oracle JRE you want to use for the Tivoli Enterprise Portal desktop client.

Browser clients

Configuration for browser clients involves the following steps.

If you are using V6.3 of the Tivoli Enterprise Portal browser client, see “Browser plug-in support for Java applets” on page 344.

- “Installing and configuring IBM Java 7” on page 348

When the browser client connects to the Tivoli Enterprise Portal Server, it downloads a Java applet. Java applets that run in a browser environment use Java plug-in technology. The applet plug-in must be registered with the browser to establish and initialize an appropriate Java runtime environment for the applet to execute in.

On Linux and UNIX, the plug-in must be manually registered with the browsers; on Windows, the Java installer automatically registers the associated Java plug-in with both Internet Explorer and Firefox. After installation, a plug-in can be de-registered from a given browser selectively using the Java control panel (see “Removing the Java plug-in on Windows” on page 345, “Removing the Java plug-in on Linux” on page 346).

- “Windows: Java 5 post-installation configuration” on page 355, “Linux: Java 5 post-installation configuration” on page 361

Depending on which Java release you are using to run the Tivoli Enterprise Portal applet, you might need to specify additional runtime parameters for the Java plug-in.

If you are using Internet Explorer under Windows, you must complete an additional step to use the Oracle JRE:

- “Identifying the version of the Oracle JRE the client should use” on page 345

Conflicts in plug-in registration can sometimes occur if both IBM and Oracle JREs are installed on Windows. You can resolve such problems by de-registering the plug-in in the browsers:

- “Removing the Java plug-in on Windows” on page 345

Tip: To open a workspace in a new tab using either Firefox or Internet Explorer, press Shift+Ctrl while selecting the workspace. This always opens a new tab whether or not your browser is set to reuse existing tabs.

Browser plug-in support for Java applets

In Tivoli Monitoring V6.3, all Java-based Tivoli Monitoring infrastructure components are being upgraded to the Java 7 runtime environment. You might have to complete additional tasks to upgrade the Tivoli Enterprise Portal component if you use the Tivoli Enterprise Portal browser or Web Start clients.

If you are doing a pristine installation of IBM Tivoli Monitoring V6.3, see Table 60.

Table 60. Tivoli Enterprise Portal client support

Java runtime environment (JRE)	Windows platforms ¹		Intel Linux platforms ¹
	Internet Explorer V8 - V9	Firefox Extended Support Release	Firefox Extended Support Release
JRE 1.6, 1.7	“Windows: Installing IBM Java 7” on page 348	“Windows: Installing IBM Java 7” on page 348	“Linux: Installing IBM Java 7” on page 357
JRE 1.5	JRE 1.5 no longer supported ²	JRE 1.5 no longer supported ²	JRE 1.5 no longer supported ²
Note: 1. Both 32-bit and 64-bit browser versions are supported. 2. Firefox Extended Support Release or higher requires Java 1.6 or higher. If you specify Java 1.5 in <code>jrelevel</code> or <code>jreDesiredLevel</code> and 1.5 is not already installed, you will be prompted to install Java 1.7 irrespective of what is specified in <code>jrelevel</code> or <code>jreDesiredLevel</code> . For Internet Explorer, the level of IBM JRE installed depends on the values specified in <code>jrelevel</code> or <code>jreDesiredLevel</code> .			

If you are upgrading from an earlier release of Tivoli Monitoring to V6.3, you already have a correctly installed and configured version of Java available on the system where the Tivoli Enterprise Portal client is used. Supported Java versions for IBM Tivoli Monitoring releases before the 6.3 release included IBM Java 5, IBM Java 6, Oracle Java 5, and Oracle Java 6. In this case, the upgrade to IBM Tivoli Monitoring V6.3 does not require any corresponding Java upgrade. You can continue to use the existing version of Java that is working successfully with your current release of the Tivoli Enterprise Portal browser and Web Start clients, after upgrading to IBM Tivoli Monitoring V6.3. No other post-installation configuration is required for the Tivoli Enterprise Portal browser or Web Start client.

Note: Starting with the IBM Tivoli Monitoring V6.3 release, use of Java 5 (1.5) is not recommended for use with the Tivoli Enterprise Portal browser client. It is strongly recommended that you upgrade to using Java 7 (1.7) that is provided with IBM Tivoli Monitoring V6.3. The use of Oracle's Java 7 is also supported for V6.3.

However, if you decide after upgrading to IBM Tivoli Monitoring V6.3 that you want to upgrade your installed version of Java to IBM Java 7, you can do so by using much of the same procedure used for a pristine installation of IBM Tivoli Monitoring V6.3. See the following sections for information about manually upgrading your existing version of Java to IBM Java 7:

- “Windows: Upgrading your existing version of Java to IBM Java 7” on page 354
- “Linux: Upgrading your existing version of Java to IBM Java 7” on page 360

Identifying the version of the Oracle JRE the client should use

If you want to use the Oracle JRE, you must identify to the Tivoli Enterprise Portal Server the version of the JRE you want the client to use.

To identify the version, update the file `jreconfig.js` found in the `\CANDLE_HOME\CNB` directory (Windows) or `CANDLE_HOME/arch/cw` (Linux/UNIX) on the computer where the portal server is installed. Assign a valid value to the following declared variable:

```
var jreDesiredLevel = "7.0"
```

The supported values are:

*	The browser client will use the default JRE for the computer (or the highest level of Java installed).
1.5.0	The browser client will use the IBM 1.5 JRE (not recommended).
1.6.0	The browser client will use the IBM 1.6 JRE.
1.7.0	The browser client will use the IBM 1.7 JRE (the default JRE for a pristine install of V6.3).
5.0	The browser client will use the latest Oracle 1.5.0_xx JRE installed.
6.0	The browser client will use the latest Oracle 1.6.0_xx JRE installed.
7.0	The browser client will use the latest Oracle 1.7.0_xx JRE installed.

Removing the Java plug-in on Windows

When either the IBM JRE or the Oracle JRE is installed on Windows, the Java installer automatically registers the associated Java plug-in with both Internet Explorer and Firefox. Plug-in registration conflicts sometimes occur on Windows when both JREs are installed.

Before you begin

The symptoms usually involve one of the following messages being displayed when you launch the Tivoli Enterprise Portal client:

"Java Plug-in detected JRE collision"

or

"Applet(s) in this HTML page requires a version of java different from the one the browser is currently using. In order to run the applet(s) in this HTML page, a new browser session is required.
Press 'Yes' to start a new browser session."

The solution is to de-register the plug-in.

About this task

To de-register the Java plug-in take the following steps:

Procedure

1. Launch the IBM Control Panel for Java or the Oracle Java Control Panel from the Windows Control Panel folder.
2. Select the **Advanced** tab.
3. Expand the branch entitled *APPLET* tag support.

There should be two entries, one for "Internet Explorer" and one for "Mozilla and Netscape".

Normally, after installation of Java, both of these boxes will be checked, meaning that the associated

Java plug-ins for those browsers have been registered. Figure 57 shows the IBM Control Panel for Java; the Oracle version is almost identical.

Note: The branch labels found in the **Advanced** tab may be different depending on whether you are using Java 1.5, 1.6, or 1.7. Figure 57 is specific to Java 1.7. The branch label for Java 1.7 is entitled **Default Java for browsers**. In addition, the actual labels associated with the browsers listed may differ between 1.5, 1.6, and 1.7. With Java 1.7, the labels are **Microsoft Internet Explorer** and **Mozilla family**.




Figure 57. Java Control Panel window

If you encounter problems loading Tivoli Enterprise Portal using Firefox, this is one of the first panels you should look at to ensure that the Java plug-in has indeed been registered. Often, a quick way to resolve any registration problems is to simply remove the check mark (if already on), press **Apply**, then check the box again, and again press **Apply**. This action will switch the option off and on so that the plug-in registration is reset and the correct Windows registry entries get re-established.

4. Remove the check mark from the box for the browser or browsers you want to unregister.

Removing the Java plug-in on Linux

Take the following steps to remove the Java plug-in:

Procedure

1. Change your current working directory to the Firefox plug-ins directory (which can differ depending on your Linux distribution):
 - For the 32 bit version of Firefox: `cd /usr/lib/mozilla/plugins/`
 - For the 64 bit version of Firefox: `cd /usr/lib64/mozilla/plugins/`
2. Depending on which Java plug-in release you currently have registered, there should be one of the following symbolic links found in the `../plugins` directory:
 - `libnpp2.so` ← next-generation plug-in provided with Java 1.6 and higher

- `libjavaplugin_oji.so` ← classic plug-in provided with Java 1.5
3. To deregister the Java plug-in, simply remove the symbolic link:
`rm -f <plug-in name>`

Java Web Start clients

No post-installation configuration of the JRE is required if the Java Web Start is used to deploy the Tivoli Enterprise Portal client.

The Tivoli Enterprise Portal client is typically deployed via Java Web Start by entering the following address in the browser's location field:

`http://teps_host:15200/tep.jnlp`

Under Windows, the last JRE installed on the machine usually controls which Java Web Start loader is associated with the Java Web Start deployment files, which have the extension `.jnlp`. If the Oracle JRE was the last JRE installed, the Oracle Java Web Start loader and associated JRE will be used for the Tivoli Enterprise Portal client. If both IBM and Oracle Java are installed on the same machine and the IBM JRE was installed last, it may be necessary to manually re-associate the `.jnlp` extension with the Oracle JRE.

The method that you use to re-associate a file extension to a program that can support that file type will vary depending on which version of Windows you have installed.

Under Linux, when you invoke Java Web Start using the Firefox browser using the URL above, you might be presented with a dialog window asking what Firefox should do with the associated `.jnlp` file (see Figure 58).

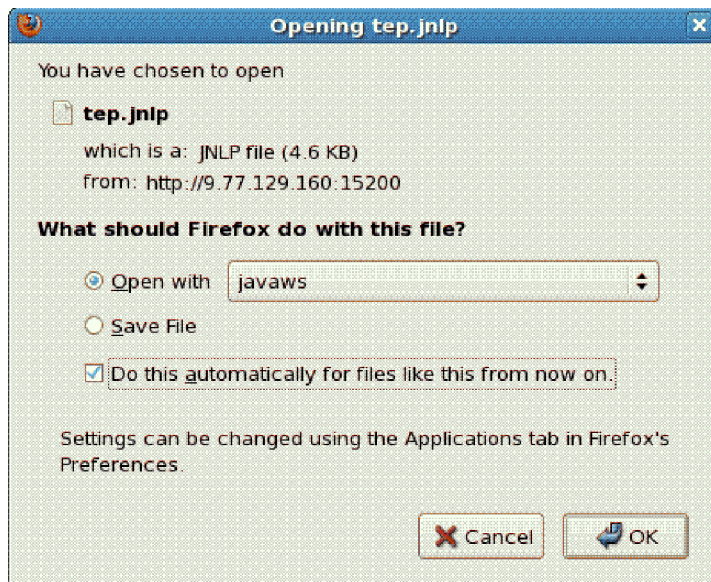


Figure 58. Launching Java Web Start client

If this happens, complete the following steps:

1. Select **Open with** and choose the Java Web Start loader program named **javaws**. This file is normally located in the `../bin` directory where you installed the JRE.
2. To ensure that this program is always used in the future for this file type, select **Do this automatically for files like this from now on**.

For more information on using and configuring Java Web Start client and setting up its environment, see "Using Web Start to download and run the desktop client" on page 367.

Installing and configuring IBM Java 7

In Tivoli Monitoring V6.3, all Java-based Tivoli Monitoring infrastructure components are being upgraded to the Java 7 runtime environment. Although this upgrade is transparent for most Tivoli Monitoring components, you might have to complete additional tasks to upgrade the Tivoli Enterprise Portal component depending on the deployment mode you use for the Tivoli Enterprise Portal.

If you use the Tivoli Enterprise Portal desktop client, the upgrade to Java 7 is transparent because the supporting Java 7 binaries and support files are embedded within the Tivoli Enterprise Portal directory structure, and are updated automatically by the Tivoli Monitoring Installer. However, if you also use the Tivoli Enterprise Portal browser or Web Start client deployments, you might have to complete some additional installation and configuration tasks to successfully upgrade the Tivoli Enterprise Portal to Java 7. This section describes the additional tasks for users of the Tivoli Enterprise Portal browser and the Web Start client.

- “Windows: Installing IBM Java 7”
- “Linux: Installing IBM Java 7” on page 357

Upgrade installation versus pristine installation

If you are upgrading from an earlier release of Tivoli Monitoring to V6.3, you already have a correctly installed and configured version of Java available on the system where the Tivoli Enterprise Portal client is used.

Supported Java versions for IBM Tivoli Monitoring releases before the V6.3 release included IBM Java 5, IBM Java 6, Oracle Java 5, and Oracle Java 6. In this case, the upgrade to IBM Tivoli Monitoring V6.3 does not require any corresponding Java upgrade. You can continue to use the existing version of Java that is working successfully with your current release of the Tivoli Enterprise Portal browser and Web Start clients, after upgrading to IBM Tivoli Monitoring V6.3. No other post-installation configuration is required for the Tivoli Enterprise Portal browser or Web Start client.

Note: Starting with the IBM Tivoli Monitoring V6.3 release, use of Java 5 (1.5) is not recommended for use with the Tivoli Enterprise Portal browser client. It is strongly recommended that you upgrade to using Java 7 (1.7) that is provided with IBM Tivoli Monitoring V6.3. The use of Oracle's Java 7 is also supported for V6.3.

However, if you decide after upgrading to IBM Tivoli Monitoring V6.3 that you want to upgrade your installed version of Java to IBM Java 7, you can do so by using much of the same procedure used for a pristine installation of IBM Tivoli Monitoring V6.3. See the following sections for information about manually upgrading your existing version of Java to IBM Java 7:

- “Windows: Upgrading your existing version of Java to IBM Java 7” on page 354
- “Linux: Upgrading your existing version of Java to IBM Java 7” on page 360

For a pristine installation of IBM Tivoli Monitoring V6.3 on Windows systems, the upgrade to IBM Java 7 is triggered when you first start the V6.3 Tivoli Enterprise Portal browser client. “Windows: Installing IBM Java 7” describes the installation of IBM Java 7.

For a pristine installation of IBM Tivoli Monitoring V6.3 on Linux systems, the installation and configuration of IBM Java 7 is recommended for use with the V6.3 Tivoli Enterprise Portal browser client. “Linux: Installing IBM Java 7” on page 357 describes the installation of IBM Java 7.

Windows: Installing IBM Java 7

The upgrade to IBM Java 7 is triggered when you first launch the V6.3 Tivoli Enterprise Portal browser client after a pristine installation.

About this task

The Java upgrade requirement is automatically detected and the graphic image and download link in Figure 59 are displayed in your browser page.




Figure 59. Java Plug-in

Complete the following steps to install IBM Java 7 for use by the Tivoli Enterprise Portal browser and Web Start clients:

Procedure

1. Click the hyperlink **Click Here to Begin**.
2. A download prompt might be displayed. Click **Run** (or similarly labelled instruction) to download and run the IBM Java 7 installer.
3. After the IBM Java 7 installer has downloaded, you might be prompted again to allow the execution of the installer to proceed. Click **Run** (or similarly labelled instruction) to begin the installation.
4. Select the default language for the installation and click **OK**.

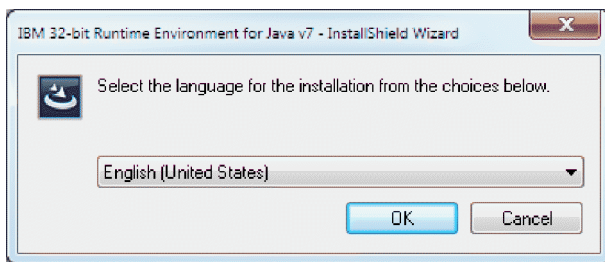


Figure 60. IBM Java 7 installation default language

5. Click **Next** on the Welcome page.




Figure 61. IBM Java 7 Welcome page

6. Click **Yes** to accept the license agreement.

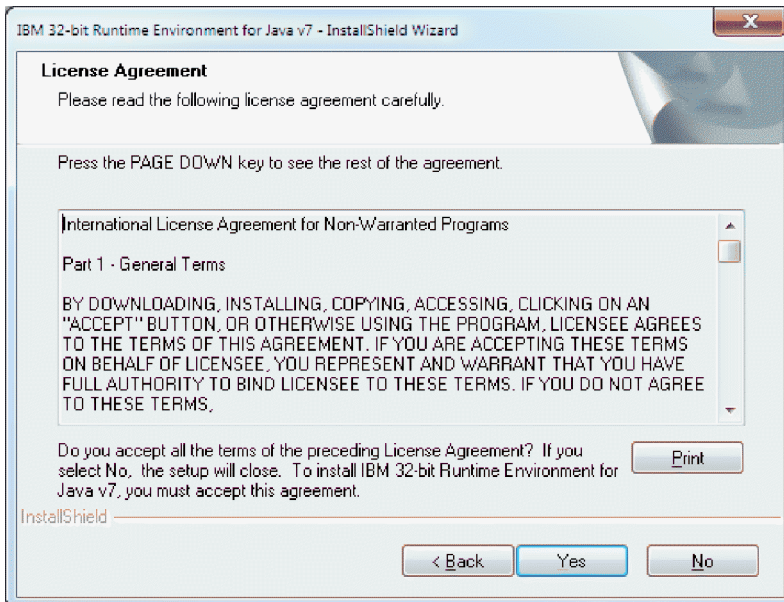


Figure 62. IBM Java 7 Software License Agreement window

7. Click **Next** to accept the default installation location for IBM Java 7, or click **Browse** to browse to a different target location on your system.




Figure 63. IBM Java 7 Choose Destination Location

8. Click **Yes** when asked to install IBM Java 7 as the System JVM.




Figure 64. Install System JVM

9. If another version of Java is already installed as the System JVM, you are asked if you want to establish IBM Java 7 as the System JVM. Click **Yes**.




Figure 65. Overwrite previous System JVM

10. A panel lists the installation options you have chosen so far. Click **Next** to begin the file copy and system update phase of the installation.

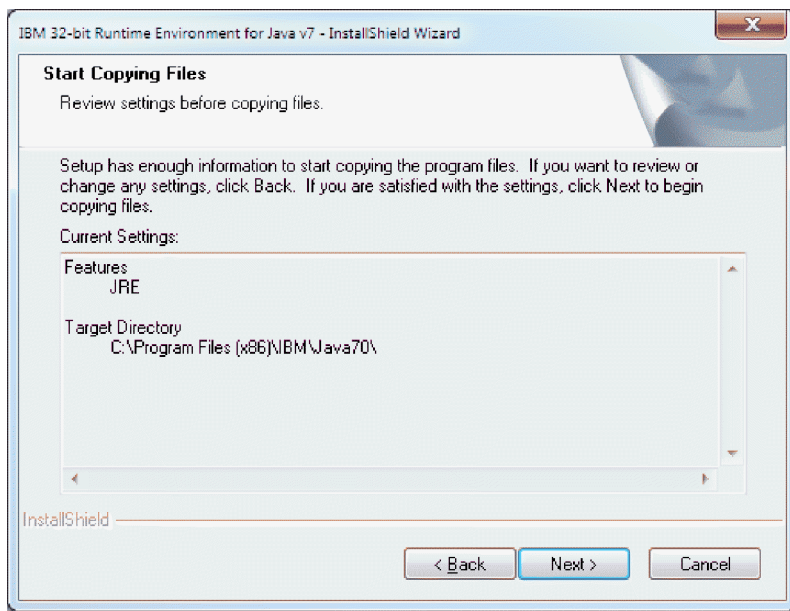


Figure 66. IBM Java 7 Start Copying Files

11. After all of the support files are copied to your machine, you are be prompted to install and register the Java Plug-in modules used by Microsoft Internet Explorer and Mozilla Firefox. Unless you are an experienced user and have specific reasons for not accepting the defaults here, click **Next** to continue.

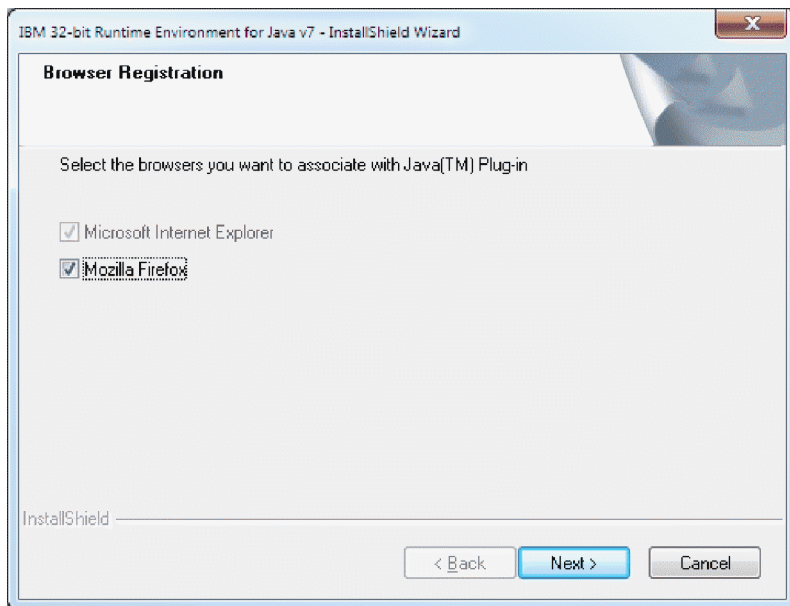


Figure 67. IBM Java 7 Browser Selection

12. After successful registration of the Java Plug-in modules, the Finish panel is displayed. Click **Finish** to complete the installation of IBM Java 7.

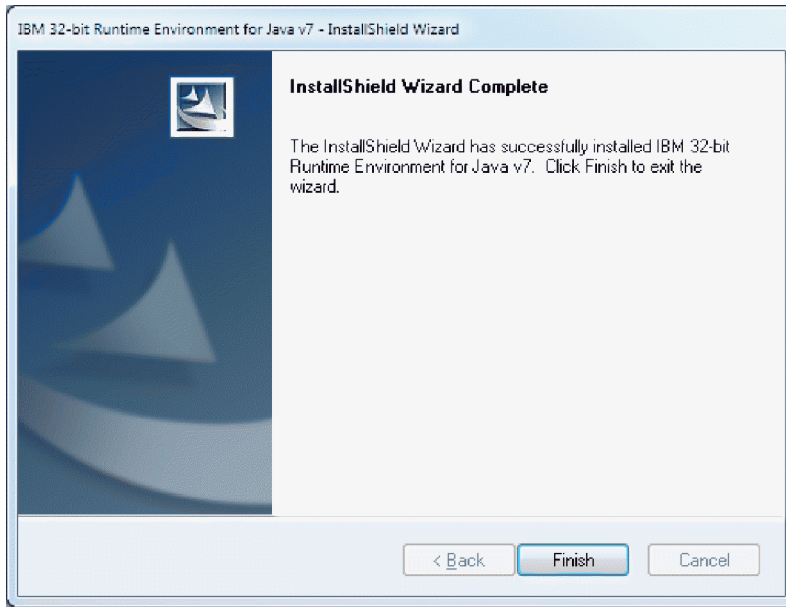


Figure 68. IBM Java 7 InstallShield Wizard Complete

13. After successful installation of Java, you can return to the window or tab in your browser containing the Tivoli Enterprise Portal client URL, and press CTRL+F5 to launch the client applet. Alternatively, you can simply click in the Address field of the browser and press the Enter key. If the client applet still fails to load, recycle the browser and then re-enter the Tivoli Enterprise Portal client URL. The applet should load successfully.

Results

You have successfully installed IBM Java 7.

Starting the Tivoli Enterprise Portal client

The first time that you start the Tivoli Enterprise Portal browser and Web Start clients, you can expect a longer than normal delay before the Tivoli Enterprise Portal login panel is displayed. The longer delay is caused by the updated IBM Tivoli Monitoring changes you make, as well as the initial population of the IBM Java 7 Java archive file (JAR) cache. Subsequent launches of the Tivoli Enterprise Portal client will be noticeably faster.

If you previously configured the Tivoli Enterprise Portal browser client to operate under IBM Java 5 or Java 6, which might be the case if you are upgrading to IBM Tivoli Monitoring V6.3 from a previous IBM Tivoli Monitoring release, you might see the following progress panel displayed during the initial launch of the client:




Figure 69. Upgrading Java applet cache

This progress panel is normal and is shown only during initial launch of the client.

Windows: Upgrading your existing version of Java to IBM Java 7

If you are upgrading to IBM Tivoli Monitoring V6.3 from a previous IBM Tivoli Monitoring release, then as stated previously in this section, no requirement exists to upgrade to IBM Java 7. Your current working Java environment for the Tivoli Enterprise Portal browser and Web Start clients will continue to work after the IBM Tivoli Monitoring upgrade is successfully completed.

About this task

If, however, you want to uplift your Java environment to IBM Java 7, you can do this using the following procedure:

Procedure

1. On the machine where the Tivoli Enterprise Portal Server is installed, locate the file named `jrelevel.js` in the `%CANDLE_HOME%\cnb` directory. If the `jrelevel.js` file is found, rename it to `jrelevel_bkup.js`.
2. On the machine where the Tivoli Enterprise Portal Server is installed, locate and open the `jreconfig.js` file for editing. The `jreconfig.js` file is located in the `<CANDLE_HOME>\cnb` directory.
3. Change the value of the `jreDesiredLevel` statement to the following value:

```
var jreDesiredLevel = "1.7.0";
```

This value references the use of IBM Java 7, and is used by the Tivoli Enterprise Portal browser client to determine which installed level of Java to use.

4. **Save** the changes you made to this file.
5. Now launch the Tivoli Enterprise Portal browser client from the machine where you normally run the Tivoli Enterprise Portal client, and follow the procedures outlined in the preceding section: “Windows: Installing IBM Java 7” on page 348.

Results

Note: Although multiple IBM Java releases can be installed and co-exist on the same machine, be aware of the following issues if you want to use multiple releases with the Tivoli Enterprise Portal browser client:

1. If you plan on switching from IBM Java 7 back to IBM Java 6 or Java 5 (or vice-versa), make sure you first update the `jreconfig.js` file located on the Tivoli Enterprise Portal Server machine as described in this section. The value to use for IBM Java 5 is the following:

```
var jreDesiredLevel = "1.5.0";
```

The value to use for IBM Java 6 is the following:

```
var jreDesiredLevel = "1.6.0";
```

In addition, the Runtime Parameters found in the Java 5 Control Panel need to be re-entered and saved. See “Windows: Java 5 post-installation configuration” on page 355 for information about configuring Java 5 for use with the Tivoli Enterprise Portal browser client.

2. You cannot launch two or more independent instances of the Tivoli Enterprise Portal browser client under Internet Explorer, from the same machine, using different IBM Java releases. If you must run more than one instance of the Tivoli Enterprise Portal client that use different IBM Java releases, then you can use another deployment mode for the second (or subsequent) instance. For example, you can run one Tivoli Enterprise Portal instance by using Internet Explorer and IBM Java 5, and a second Tivoli Enterprise Portal instance by using Web Start or desktop deployments running IBM Java 6.

Windows: Java 5 post-installation configuration

After installing IBM Java 5 you must configure the JVM to successfully launch the Tivoli Enterprise Portal browser and Web Start clients.

About this task

Complete the following steps to configure IBM Java 5 for Windows:

Procedure

1. Open your Windows **Control Panel** folder and locate the IBM Java 5 application icon.



Figure 70. IBM Java 5 application icon

2. To launch the Java Control Panel, double-click the IBM Java 5 application icon.
3. On the Java Control Panel, click the Java tab.




Figure 71. Java Control Panel

4. Click **View** to list the installed versions of the Java runtime and to configure the Java runtime arguments for the IBM Java 5 environment.
5. The Java 1.5 platform row should be the first (or only) row in the Java Runtime Environment list. Click in the **Runtime Parameters** entry field for this row and type (or copy-and-paste) the following argument string into the parameters field:
`-Xms128m -Xmx256m -Xverify:none`




Figure 72. Java Runtime Environment list

6. Press the Enter key to commit the changes to the parameters field, and click **OK** to close the JRE Settings sub-panel.
7. Click **Apply** to save your changes before continuing with the rest of the configuration.
8. Click the **Advanced** tab and expand the branches in the property tree, as shown in Figure 73.

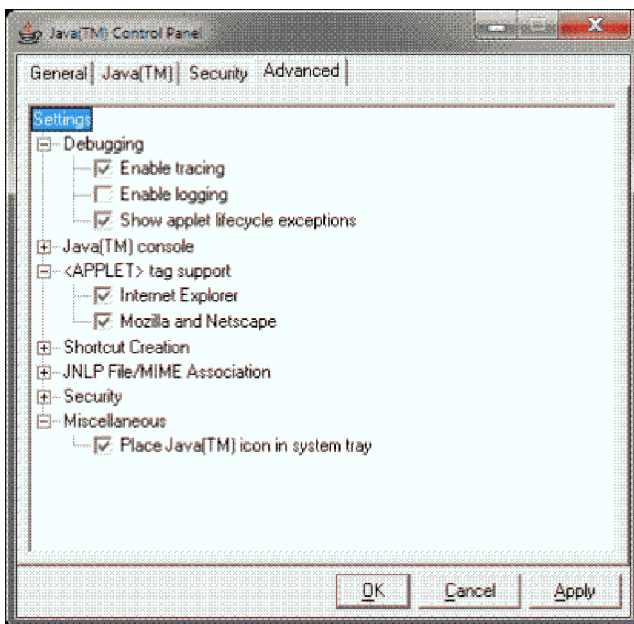


Figure 73. Java Control Panel Advanced tab

Configure the following properties in the **Advanced** tab:

- a. Under the **Debugging** branch, select the following two properties:
 - Enable tracing
 - Show applet lifecycle exceptions
- b. Under the **<APPLET> tag support** branch, select the following two properties (if not already selected):
 - Microsoft Internet Explorer
 - Mozilla and Netscape

Note: Only one of these properties might be listed, depending on what browsers you have installed on the machine.

- c. Under the **Miscellaneous** branch, select the following property:
 - **Place Java (TM) icon in system tray**

Click **OK** to continue.

9. Click **Apply** to commit the property changes, and click **OK** to save and close the Java Control Panel.
10. You must recycle the browser for some of the configuration changes to take effect. Close the browser that you are using to launch the Tivoli Enterprise Portal browser or Web Start client.

Results

You have successfully configured IBM Java 5.

Linux: Installing IBM Java 7

There are two alternative methods for installing IBM Java 7.

“Method 1: installation of an rpm-based package” and “Method 2: retrieval and extraction of a Linux .tar gzip archive” on page 359. This allows you to choose between an rpm-based installation, or a simple download and extract operation to a location of your choice.

Method 1: installation of an rpm-based package

The IBM Java 7 installation package is provided as an .rpm package file located on the system where the Tivoli Enterprise Portal Server component is installed. Support for the Tivoli Enterprise Portal browser client on Linux has been tested using the Mozilla Firefox browser, releases 3.6 through the most recent Extended Support Release at the 32-bit and 64-bit level.

About this task

The upgrade to IBM Java 7 is triggered when you first launch the V6.3 Tivoli Enterprise Portal browser client after a pristine installation. The Java upgrade requirement is automatically detected and the graphic image and download link in are displayed in your browser page.




Figure 74. Java Plug-in is not installed

Complete the following steps to install IBM Java 7 for use with the Tivoli Enterprise Portal browser and Web Start clients in a Linux environment:

Procedure

1. Click the hyperlink **Click Here to Begin**.
2. A download prompt might be displayed. Click **OK** (or similarly labelled instruction) to download and process the package.




Figure 75. Download .rpm package file

3. Depending on the authority level associated with your user credentials, you might be prompted for Administrative privileges before the installation of IBM Java 7 can continue. If required, provide your root password when prompted.
4. A final prompt to proceed with the installation of the IBM Java 7 package might be displayed. Click Apply (or similarly labelled instruction) to continue.




Figure 76. IBM Java 7 installing packages

5. After the IBM Java 7 package has been successfully installed, you might receive a completion message. Click **OK** to continue.




Figure 77. IBM Java 7 successfully installed

6. After the package manager has installed the IBM Java 7 package, you must register the Java plug-in with your Firefox browser. To register the Java plug-in, you must symbolically link it to the plug-in directory associated with your browser.

Change your working directory to the Firefox plug-ins directory (which can differ depending on your Linux distribution and whether you are using the 32 bit or 64 bit version of Firefox):

- For the 32 bit version of Firefox: `cd /usr/lib/mozilla/plugins/`
 - For the 64 bit version of Firefox: `cd /usr/lib64/mozilla/plugins/`
7. Now create a symbolic link to the plug-in using the following command:
 - For the 32 bit Java plug-in for IBM Java 7: `ln -s /opt/ibm/java-i386-70/jre/lib/i386/libnpjp2.so .`
 - For the 64 bit Java plug-in for IBM Java 7: `ln -s /opt/ibm/java-x86_64-70/jre/lib/amd64/libnpjp2.so .`
 8. After successful installation of Java and registration of the Java plug-in, you can return to the window or tab in your browser containing the Tivoli Enterprise Portal client URL, and press CTRL+F5 to launch the client applet. Alternatively, you can click in the Address field of the browser and press the Enter key. If the client applet still fails to load, recycle the browser and then re-enter the Tivoli Enterprise Portal client URL. The applet should load successfully.



Method 2: retrieval and extraction of a Linux .tar gzip archive

Complete the following steps to install IBM Java 7 for use with the Tivoli Enterprise Portal browser and Web Start clients in a Linux environment.

Procedure

1. Download the IBM Java 7 tgz file from the machine where your Tivoli Enterprise Portal Server is located:
 - The 32 bit archive file for IBM Java 7 is named `ibm-java7.tgz`. The 32 bit archive file is available from the following location:
`http://teps_hostname:15200/java/ibm-java7.tgz`
 - The 64 bit archive file for IBM Java 7 is named `ibm-java7_64.tgz`. The 64 bit archive file is available from the following location:
`http://teps_hostname:15200/java/ibm-java7_64.tgz`

Alternatively, if you prefer to use a different method for retrieving this file, the following directory location is where you can find the file on the Tivoli Enterprise Portal Server machine:

-  `<itm_install_directory>\cnb\java`
-  `<itm_install_dir>/<arch>/cw/java`

Save the file in a directory of your choice, for example `/home`.

2. As the **root** user, open a shell prompt.
3. Create a directory where you plan on installing IBM Java 7 on your client machine. A suggested directory location is `/opt/ibm`.
4. Extract the IBM Java 7 installation archive contents using the following command (here it is assumed the `ibm-java7_64.tgz` file was downloaded to the `/home` directory and will be installed in the `/opt/ibm` directory):
`tar -xvzf /home/ibm-java7_64.tgz -C /opt/ibm`
5. After you have extracted the IBM Java 7 installation archive contents you must register the Java plug-in with your Firefox browser. To register the Java plug-in, you must symbolically link it to the plug-in directory associated with your browser. See steps 6 on page 358, 7, and 8 in method 1 for information on how to complete this task.

Starting the Tivoli Enterprise Portal client

The first time that you start The Tivoli Enterprise Portal browser and Web Start clients, you can expect a longer than normal delay before the Tivoli Enterprise Portal login panel is displayed. The longer delay is caused by the updated IBM Tivoli Monitoring changes you make, as well as the initial population of the IBM Java 6 Java archive file (JAR) cache. Subsequent launches of the Tivoli Enterprise Portal client will be noticeably faster.

If you previously configured the Tivoli Enterprise Portal browser client to operate under IBM Java 5, as might be the case if you are upgrading to IBM Tivoli Monitoring V6.3 from a previous IBM Tivoli Monitoring release, you might see the following progress panel displayed during the initial launch of the client:




Figure 78. Upgrading Java applet cache

This progress panel is normal and is shown only during initial launch of the client.

Linux: Upgrading your existing version of Java to IBM Java 7

If you are upgrading to IBM Tivoli Monitoring V6.3 from a previous IBM Tivoli Monitoring release, then as stated previously in this section, no requirement exists to upgrade to IBM Java 7. Your current working Java environment for the Tivoli Enterprise Portal browser and Web Start clients will continue to work after the IBM Tivoli Monitoring upgrade is successfully completed.

About this task

You can also complete the following steps if you want to upgrade your Java environment to IBM Java 7 for the browser client:

Procedure

1. On the machine where the Tivoli Enterprise Portal Server is installed, locate the file named `jrelevel.js` in the following directory:

- **Windows** <CANDLE_HOME>\cnb
- **Linux** **UNIX** <CANDLE_HOME>/cw

If the `jrelevel.js` file is found, rename it to `jrelevel_bkup.js`.

2. On the machine where the Tivoli Enterprise Portal Server is installed, locate and open the `jreconfig.js` file for editing. The `jreconfig.js` file is located in the following directory:

- **Windows** <CANDLE_HOME>\cnb
- **Linux** **UNIX** <CANDLE_HOME>/cw

3. Change the value of the `jreDesiredLevel` statement to the following value:

```
var jreDesiredLevel = "1.7.0";
```

This value references the use of IBM Java 7, and is used by the Tivoli Enterprise Portal browser client to determine which installed level of Java to use.

4. Save the changes you made to this file.
5. Start the Tivoli Enterprise Portal browser client from the machine where you normally run the Tivoli Enterprise Portal client, and follow the procedures outlined in the preceding sections: “Linux: Installing IBM Java 7” on page 357.

Results

Note: Although multiple IBM Java releases can be installed and co-exist on the same machine, be aware of the following issues if you want to use both of these releases with the Tivoli Enterprise Portal browser client:

1. If you plan on switching from IBM Java 7 back to IBM Java 6 or Java 5 (or vice-versa), make sure you first update the `jreconfig.js` file located on the Tivoli Enterprise Portal Server machine as described in this section. The value to use for IBM Java 6 is the following:

```
var jreDesiredLevel    = "1.6.0";
```
2. When switching between IBM Java releases, you must re-enter and save and Runtime Parameters you many have entered previously for the appropriate JRE. See “Linux: Java 5 post-installation configuration” for information about configuring Java 5 for use with the Tivoli Enterprise Portal browser client.
3. The version of the Java plug-in used by Firefox is based on the registration of the plug-in object. Follow steps 6 on page 358 and 7 on page 359 in “Linux: Installing IBM Java 7” on page 357 to establish a symbolic link to the plug-in object, depending on what version of Java you are switching over to use (Java 5 or Java 6). For Java 5, the plug-in object is named `libjavaplugin_oji.so`. For Java 7, the plug-in object is named `libnpjp2.so`.

Linux: Java 5 post-installation configuration

After installing IBM Java 5 you must configure the JVM to successfully launch the Tivoli Enterprise Portal browser and Web Start clients.

About this task

Complete the following steps to configure IBM Java 5 for Linux:

Procedure

1. Open a shell prompt, making sure you do so as the **root** user.
2. Launch the IBM Java 5 Control Panel application:

```
cd /opt/ibm/java-i386-50/jre/bin
./ControlPanel
```
3. On the Java Control Panel, click the **Java** tab.

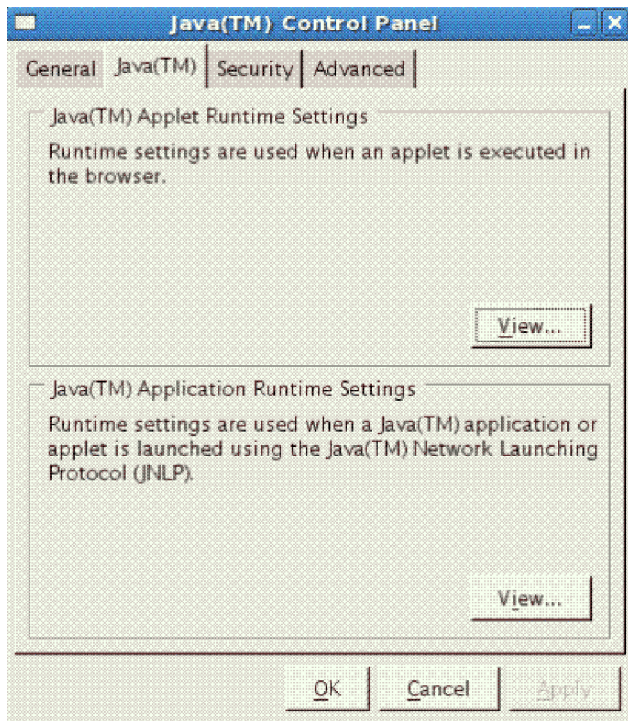


Figure 79. Java Control Panel

4. On the Java tab, click **View** to see the installed versions of the Java runtime, and to configure the Java runtime arguments for the IBM Java 5 environment.
5. Click in the **Runtime Parameters** entry field associated with the Java 1.5 platform row, choose the one associated with IBM Java Path, and type (or copy-and-paste) the following argument string into the parameters field:
`-Xms128m -Xmx256m -Xverify:none`

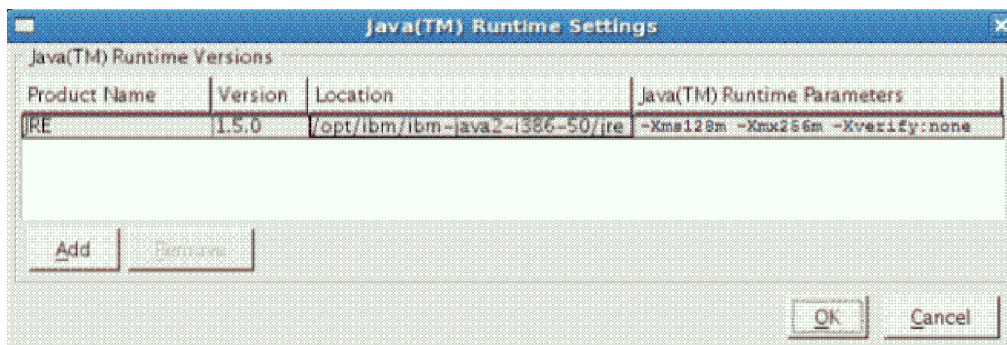


Figure 80. Java Runtime Environment Settings

6. Press the **Enter** key to commit the changes to the parameters field, and click **OK** to close the JRE Settings sub-panel.
7. The Java Control Panel is displayed, click **Apply** to save your changes before continuing with the rest of the configuration.
8. Click the **Advanced** tab and expand the branches in the property tree, as shown in Figure 81 on page 363.




Figure 81. Java Control Panel

Configure the following properties in the **Advanced** tab:

- a. Under the **Debugging** branch, select the following two properties:
 - Enable tracing
 - Show applet lifecycle exceptions
 - b. Under the **Command to launch default browser** branch, make sure you have the correct path for launching the Firefox browser.
9. Click **Apply** to commit the property changes, and click **OK** to save and close the Java Control Panel.
 10. You must recycle the Firefox browser for some of the configuration changes to take effect.

Specifying the browser used for online help

If you are running the desktop client on Linux, or you want to view the online help with some browser other than Internet Explorer on Windows, you must specify to the portal server the location of the browser you want to use.

- “Windows: Specifying the browser location”
- “UNIX and Linux: Specifying the browser location” on page 364
- “Web Start: Specifying the browser location” on page 365

Windows: Specifying the browser location

Use the Manage Tivoli Enterprise Monitoring Services utility to change the location of the browser that the browser or desktop client uses.

Procedure

1. Launch Manage Tivoli Enterprise Monitoring Services (**Start → (All) Programs → IBM Tivoli Monitoring → Manage Tivoli Monitoring Services**).
2. In the Manage Tivoli Enterprise Monitoring Services window, right-click the browser or desktop client and select **Reconfigure**.

The Configure the Tivoli Enterprise Portal Browser window is displayed. (If you are configuring the desktop client, the Configure Application Instance window is displayed.)

3. Scroll down in the list of variables until you see the `kjr.browser.default` variable.
4. Double-click `kjr.browser.default`.

The Edit Tivoli Enterprise Portal Browser Parm window is displayed.

5. In the Value field, type the path and the application name of the alternative browser application. For example:
`C:\Program Files\Mozilla Firefox\firefox.exe`
6. Check the **In Use** box.
7. Click **OK** to close the editing window and save the change.
8. Click **OK** to close the reconfiguration window.

UNIX and Linux: Specifying the browser location

To change a property such as the location of the Web browser that the Tivoli Enterprise Portal browser client launches in UNIX, update the shell script file or files that are run and the template that is used when the browser client is configured to create the script file or files that are run.

You might have to update one or more of the files list in Table 61:

Note: All file paths are relative to your *install_dir* directory where you installed IBM Tivoli Monitoring.

Table 61. File locations for changing application properties for UNIX and Linux

File location	Purpose of file
<code>bin/cnp.sh</code>	The default shell script that launches the Tivoli Enterprise Portal browser client.
<code>bin/cnp_instance.sh</code>	The shell script for a specific instance you have created, where <i>instance</i> is the name of the instance that launches the Tivoli Enterprise Portal browser client.
<code>platform/cj/original/cnp.sh_template</code>	<p>The template from which the <code>bin/cnp.sh</code> and <code>bin/cnp_instance.sh</code> shell scripts are generated during configuration, where <i>platform</i> is the code for the operating system platform on which IBM Tivoli Monitoring is installed. For example: <i>li6243</i> for Linux 2.4 on a 32-bit Intel CPU).</p> <p>If you change only <code>bin/cnp.sh</code> or <code>bin/cnp_instance.sh</code> and do not change this template, the next time you configure the client, a new version of the script is created without the changes you made to <code>bin/cnp.sh</code> or <code>bin/cnp_instance.sh</code>.</p>

To change the location of the Web browser you must change the above file or files to include a new property:

1. Go to the `install_dir/bin/cnp.sh` and edit the `cnp.sh` shell script.
2. Add your Web browser location to the last line of the file. In the example below, the Web browser location is `/opt/foo/bin/launcher`. `-Dkjr.browser.default=/opt/foo/bin/launcher`

Important: The line is very long and has various options on it, including several other `-D` options to define other properties. It is very important to add the option in the correct place.

If the last line of your `bin/cnp.sh` originally looked like the following:

```

${JAVA_HOME}/bin/java -showversion -noverify -classpath ${CLASSPATH}
-Dkjr.trace.mode=LOCAL -Dkjr.trace.file=/opt/IBM/ITM/logs/kcjrsl.log
-Dkjr.trace.params=ERROR -DORBTcpNoDelay=true -Dcnp.http.url.host=
-Dvbroker.agent.enableLocator=false
-Dhttp.proxyHost=
-Dhttp.proxyPort=candle.fw.pres.CMWApplet 2>& 1 >> ${LOGFILENAME}.log

```

To set the browser location to */opt/foo/bin/launcher*, change the line to look like the following:

```

${JAVA_HOME}/bin/java -showversion -noverify -classpath ${CLASSPATH}
-Dkjr.browser.default=/opt/foo/bin/launcher
-Dkjr.trace.mode=LOCAL -Dkjr.trace.file=/opt/IBM/ITM/logs/kcjrsl.log
-Dkjr.trace.params=ERROR -DORBTcpNoDelay=true -Dcnp.http.url.host=
-Dvbroker.agent.enableLocator=false
-Dhttp.proxyHost=
-Dhttp.proxyPort=candle.fw.pres.CMWApplet 2>& 1 >> ${LOGFILENAME}.log

```

Web Start: Specifying the browser location

Java Web Start deployed applications are described in *.jnlp* deployment files. For IBM Tivoli Monitoring, there is one deployment file that describes the core Tivoli Enterprise Portal framework component and associated JAR files, and one deployment file for each and every Tivoli Enterprise Portal-based monitoring solution that is installed.

The core Tivoli Enterprise Portal Server deployment file is named *tep.jnlp*. The application deployment file is typically called *kxx_resources.jnlp* or *kxx.jnlp*, where *xx* is the application identifier (a product code, such as **nt**, **ux**, or **lz**).

- On a Windows computer where the Tivoli Enterprise Portal Server is installed, the file is located in *itminstall_dir\CNB* (for example, *c:\IBM\ITM\CNB*).
- On a Linux computer where the Tivoli Enterprise Portal Server is installed, the file is located in *itminstall_dir/arch/cw* (for example, */opt/IBM/ITM/1x8266/cw*).

The deployment file instances are generated whenever the Tivoli Enterprise Portal Server is installed or reconfigured (for example, when adding a new monitoring solution to the environment). The contents of these files are based upon two template deployment files (*.jnlp*pt). The core Tivoli Enterprise Portal template deployment file is called *tep.jnlp*pt. The application template deployment file is named *component.jnlp*pt.

- On a Windows computer where the Tivoli Enterprise Portal is installed, the file is located in *itminstall_dir\Config* (for example: *c:\IBM\ITM\Config*).
 - On a UNIX computer where the Tivoli Enterprise Portal is installed, the file is located in *itminstall_dir/config* (for example, */opt/IBM/ITM/config*).
1. To add or modify JVM arguments (such as maximum heap size) or other Tivoli Enterprise Portal-based properties (such as RAS1 trace options), you must edit either the *tep.jnlp* deployment file or the *tep.jnlp*pt deployment template file. The deployment file is nothing more than XML syntax that describes the Web Start application being deployed. The **<resources>** element is used to define the JVM arguments, the Tivoli Enterprise Portal properties, JAR files, and references to component deployment files.
 - Modify the *tep.jnlp* file if the change will be temporary (for example, setting a trace option for gathering further diagnostics).
 - Modify the *tep.jnlp*pt template file if the change will be long-term (for example, increasing the maximum heap size to accommodate a larger monitored environment or increased event load).

If you modify the deployment template file, make sure you then reconfigure the Tivoli Enterprise Portal Server in order to regenerate the instance-level *.jnlp* deployment files with your changes.

Note: Changes to `tep.jnlpt` are not preserved when you migrate to a new release. A backup of the `tep.jnlpt` file will be made during migration; you must manually compare this backup to the new version installed during migration to determine what changes need to be manually incorporated into the new `tep.jnlpt`.

2. To specify the location of the browser to use to display the online help, add the following property to the **<resources>** section of the appropriate file:

```
<property name="kjr.browser.default" value="fully-qualified path of browser executable">
```

Windows example:

```
<resources os="Windows">
  <jar href="classes/browser-winnt.jar"/>
  <jar href="classes/browser-core-winnt.jar"/>
  <property name="kjr.browser.default" value="C:\Program Files\Internet Explorer\iexplore.exe"/>
</resources>
```

Linux example:

```
<resources os="Linux">
  <jar href="classes/browser-li.jar"/>
  <jar href="classes/browser-core-li.jar"/>
  <property name="kjr.browser.default" value="/usr/bin/firefox"/>
</resources>
```

Note:

For the online help to display on a Linux or Unix computer, the `kjr.browser.default` value must specify Firefox as the default browser.

`kjr.browser.default` is not the only property you can specify using the **<property>** keyword. You can include any client parameters that are specific to your operating system.

Starting the Tivoli Enterprise Portal client

After you have successfully installed and configured all the components of your IBM Tivoli Monitoring environment, you can verify the installation and configuration by launching the Tivoli Enterprise Portal to view monitoring data.

You can access the Tivoli Enterprise Portal using either the desktop client or the browser client.

Your monitoring server and portal server must be running for the portal client to start successfully.

Starting the desktop client

Follow these steps to start the desktop client.

About this task

Windows

1. Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Tivoli Enterprise Portal**.
2. Type your user ID and password in the logon window. The default user ID is **sysadmin**.
3. Click **OK**.

Linux

Run the following command to start the portal desktop client:

```
./itmcmd agent start cj
```

Starting the browser client

Follow these steps to start the browser client:

Procedure

1. Start the browser.
2. Type the URL for the Tivoli Enterprise Portal into the **Address** field of the browser:
`http://systemname:15200/`
where the *systemname* is the host name of the computer where the Tivoli Enterprise Portal Server is installed, and 15200 is the port number for the browser client. 15200 is the default port number for the browser client. Your portal server might have a different port number assigned.
3. Click **Yes** on the Warning - Security window.
4. Type your user ID and password in the logon window. The default user ID is **sysadmin**.
5. Click **OK**.

Using Web Start to download and run the desktop client

A desktop client obtained from the Tivoli Enterprise Portal Server through IBM Web Start for Java benefits from centralized administration from the server. Like the browser client, it is automatically configured with the latest updates each time you start the client, and there is no need to configure application support.

Before you use IBM Web Start for Java to download the desktop client from the Tivoli Enterprise Portal Server:

- The Tivoli Enterprise Portal Server must be installed. (See “Installing the Tivoli Enterprise Portal Server” on page 226.)
- IBM 32 bit or 64 bit Runtime Environment for Windows, version 7.0 must be installed on the computer to which you want to download the desktop client.

You can download the IBM JRE installer from the Tivoli Enterprise Portal Server (see “Browser plug-in support for Java applets” on page 344). The IBM JRE must be installed as the system JVM.

If you want to run the desktop client on a system that already has a Tivoli Management Services base component installed (such as a monitoring server or the portal server), there is no need to install the IBM JRE. The correct version of the IBM JRE is installed with the Tivoli Management Services component.

If you run the desktop client using Web Start instead of installing it from the installation media, you must configure the JRE to enable tracing for the desktop client (see “Enabling tracing for the JRE”).

Enabling tracing for the JRE

Log files are not created for the desktop client launched through Web Start unless you enable tracing for the JRE.

About this task

The logs for a desktop client run using Web Start are located in a different place than logs for the browser client and the desktop client installed from the media. On Windows computers, the logs for the Web Start client are located in the %USERPROFILE%\AppData\LocalLow\IBM\Java\Deployment\log directory. On Linux computers, the logs are located in the .java/deployment directory of the home directory of the user ID under which the Java JRE was installed. Java Web Start will create a uniquely named trace file for every independent launch of the application. The files are named **javawsnnnnnn.trace**, where *nnnnnn* is an arbitrary five-digit identifier.

Complete the following steps to enable tracing:

Procedure

1. Launch the IBM Control Panel for Java.
 - **Windows** Select **Start → Control Panel**, then double-click IBM Control Panel for Java.
You must switch to the Classic view to see and select the Control Panel.
 - **Linux** Change to *install_dir/jre/platform/bin* and run Control Panel:
`./ControlPanel`
2. Select the **Advanced** tab.
3. Expand the Debugging node in the **Settings** tree and check **Enable Tracing**.
4. Click **OK** to save the setting and close the Java Control Panel.

Downloading and running the desktop client

You can use any of the following three methods to download and run the desktop client using Web Start.

- Entering the URL of the portal server in a browser
- Launching the client from the IBM Java Control Panel
- Entering a URL from the command-line

The first time you launch the desktop client, you are prompted to create a shortcut. After that, you can launch the client using the shortcut.

Entering the URL of the portal server in a browser

Use the following procedure to launch the desktop client from a browser:

1. Start the browser on the computer on which you want to use the desktop client.
2. Enter the following URL in the **Address** field of the browser:
`http://systemname:15200/`
where *systemname* is the fully qualified host name of the computer where the Tivoli Enterprise Portal Server is installed (for example, `myteps.itmlab.company.com`).
3. Click **Run** on the security message.
4. You are asked if you want to create a shortcut on your desktop for the Tivoli Enterprise Portal. Click **Yes** if you want to create the shortcut.
The desktop client starts and displays the logon window.

Note: If IBM Java is not the system JVM, you cannot use this shortcut. You must create your own. See “Manually creating a shortcut for the Web Start client” on page 369.

5. Enter the user ID and password to log on to the Tivoli Enterprise Portal or click **Cancel** if you do not want to log on at this time. (The default user ID is **sysadmin**.)


Note: If you set the RAS trace option for the Tivoli Enterprise Portal client as documented in , when you recycle the client the `kcjras1.log` should be created in the location from which the client was launched. On Windows this defaults to `\Documents and Settings\userid\Desktop`.

Launching the desktop client from the IBM Java Control Panel

IBM Java provides a control panel for managing both Java applets deployed via the Java plug-in, and Java applications deployed via Web Start. Complete the following steps to launch the desktop client using the control panel:

1. Launch the IBM Java Control Panel:
 - **Windows** In the Windows control panel, double-click **IBM Java Control Panel**.

Note: You must be in the Classic view to see **IBM Java Control Panel**.

-  Change to *install_dir/jre/platform/bin* directory (the default directory is */opt/IBM/ITM/jre/platform/bin* and enter *./ControlPanel*).
2. On the **General** tab, in the Temporary Internet Files section, click **Settings**. The Temporary Files Settings window is displayed.
 3. Click **View Applications**.
 4. On the **User** tab, select Tivoli Enterprise Portal, and then click **Launch Online**.
Web Start downloads and starts the desktop client. When the application is launched, you can close the Control Panel windows.

Launching the desktop client from the command-line

Complete the following steps to launch the desktop client using Web Start from the command-line:

1. Open a command prompt and change to the directory where Web Start is installed.
On Windows, the default directory is *C:\Program Files\IBM\Java70\jre\bin*. On Linux, the default directory is *install_dir/jre/platform/bin*.
2. Enter the following command:

```
javaws http://systemname:15200/tep.jnlp (Windows)
./javaws http://systemname:15200/tep.jnlp (Linux)
```

 where *systemname* is the fully qualified host name of the computer where the Tivoli Enterprise Portal Server is installed (for example, *myteps.itmlab.company.com*).

Web Start downloads and launches the desktop client.

Manually creating a shortcut for the Web Start client

On Windows, the Web Start executable file for the default Java JVM is copied to a Windows system directory. When you let Web Start create a short cut for launching the desktop client, it uses the file in the system directory as the target. You can create a shortcut manually.

About this task

To create a shortcut to launch the desktop client using Web Start:

Procedure

1. Right-click on the Windows desktop and select **New → Shortcut** from the popup menu.
2. In the Create Shortcut window, type the following path or click **Browse** and navigate to the executable as shown.
 - For the IBM 32-bit JRE on a Windows 32-bit operating system:
C:\Program Files\IBM\Java70\jre\bin\javaws.exe http://systemname:15200/tep.jnlp
 - For the IBM 32-bit JRE on a Windows 64-bit operating system:
C:\Program Files (x86)\IBM\Java70\jre\bin\javaws.exe http://systemname:15200/tep.jnlp
 - For the IBM 64-bit JRE on a Windows 64-bit operating system:
C:\Program Files\IBM\Java70\jre\bin\javaws.exe http://systemname:15200/tep.jnlp
3. Click **Next** and type a name for the shortcut in the Select a Title for the Program window. For example:
ITM Web Start client
4. Click **Finish**.
The shortcut appears on your desktop.

Installing product maintenance

The installation procedures in this chapter contain instructions for new installations. Follow the same procedures for upgrading or updating an existing installation. Fix packs (that is, product maintenance) use the same installer as a pristine installation except that in this case, the installer runs only on a system that is already installed and configured. This is true when installing both generally available fix packs and interim fixes.

An *upgrade* is an installation that replaces a previous release or fix pack level of the product or component with a later release or fix pack level. An *update* is a modification to an existing installation at the same release or fix pack level.

Note that, when you upgrade or update an installation, configuration windows for components that are already configured might not be displayed. Skip those steps that do not apply; there is no need to rerun any configuration step unless specifically called for in the fix pack's documentation.

The instructions sometimes reference the default path for the installation of IBM Tivoli Monitoring. If you did not install the product in the default path, you must specify the correct path whenever the upgrade procedures require a path specification.

Chapter 11. Deploying monitoring agents across your environment

IBM Tivoli Monitoring provides the ability to deploy monitoring agents from a central location, the monitoring server.

Just as there are two types of monitoring agents, there are two types of agent deployment:

- OS agent deployment from the installation image or using the **tacmd createNode** command
- Non-OS agent (such as the DB2 for Linux, UNIX, and Windows agent) deployment using the Tivoli Enterprise Portal GUI (for other non-OS agents) or the **tacmd addSystem** command

Table 62 describes the steps required to set up and manage remote agent deployment:

Table 62. Remote agent deployment tasks

Goal	Where to find this information
Create and populate the agent deploy depot with installable agent images.	"Populating your agent depot"
View and change the contents of the agent depot.	"Managing your agent depot" on page 374
Use one agent depot for all the monitoring servers in your monitoring environment.	"Sharing an agent depot across your environment" on page 375
Use prerequisite scanning before or during agent deployment	"Remote prerequisite scanning" on page 51
Deploy an OS agent.	"Deploying OS agents" on page 375
Deploy a non-OS agent.	"Deploying non-OS agents" on page 377
Deploy a group of agents simultaneously.	"Bulk agent deployment" on page 382

You can also use the remote agent deployment function to configure deployed agents and install maintenance on your agents. For information, see the *IBM Tivoli Monitoring Administrator's Guide*. See the *IBM Tivoli Monitoring Command Reference* for commands that you can use to perform these tasks.

Important: Run the **tacmd login** command before executing commands from the tacmd library. This requirement does not apply to the **addBundles** command. Run the **tacmd logoff** command after you finish using the tacmd command library."

Populating your agent depot

The *agent depot* is an installation directory on the monitoring server from which you deploy agents and maintenance packages across your environment. Before you can deploy any agents from a monitoring server, you must first populate the agent depot with bundles. A *bundle* is the agent installation image and any prerequisites.

When you add a bundle to the agent depot, you need to add the bundle that supports the operating system to which you want to deploy the bundle. For example, if you want to deploy a DB2 for Linux, UNIX, and Windows agent bundle to a computer running HP-UX, add the HP-UX-specific agent bundle to the depot. If your depot directory is on Windows and you want to deploy the DB2 for Linux, UNIX, and Windows agent to HP-UX, load the HP-UX bundle from the DB2 for Linux, UNIX, and Windows agent installation media for HP-UX. (If you are installing from different media for each platform type, for example, Windows, AIX and Solaris, HP-UX, Linux, you need to add the bundle from the specific platform media for the component.)

You can have an agent depot on each monitoring server in your environment or share an agent depot, as described in “Sharing an agent depot across your environment” on page 375. If you choose to have an agent depot for each monitoring server, you can customize the agent depot based on the types of bundles that you want to deploy and manage from that monitoring server. For example, if you have a monitoring server dedicated to monitoring the DB2 for Linux, UNIX and Windows agents in your environment, populate the depot with DB2-related agent bundles. If you deploy an agent from a remote monitoring server, you must have a agent bundle in the depot available to the monitoring server.

Note: Agent depots cannot be located on a z/OS monitoring server.

There are two methods to populate the agent depot:

- “Populating the agent depot from the installation image”
- “Populating the agent depot with the tacmd addBundles command” on page 373

Populating the agent depot from the installation image

You can use the installation image to populate the agent depot only when you are populating the depot with bundles for the same operating system as your monitoring server.

Use the following sections to populate your agent depot from the installation image:

- “Windows: Populating the agent depot during installation”
- “Linux and UNIX: Populating the agent depot during installation”

You can use the installation image to add a bundle for a Windows agent to a Windows monitoring server, but you cannot use the Linux installation image to add a Linux bundle to a Windows monitoring server. If you need to add bundles for operating systems other than that used by your monitoring server, use the **tacmd addBundles** command, as described in “Populating the agent depot with the tacmd addBundles command” on page 373.

Attention: Load only Tivoli-provided product agent bundles into the IBM Tivoli Monitoring deployment depot. User-provided or customized bundles are not supported. Use only Tivoli-provided tacmd commands to process bundles and to execute agent deployments. Manual manipulation of the depot directory structure or the bundles and files within it is not supported and may void your warranty.

Windows: Populating the agent depot during installation

Follow the steps in this section to populate the agent depot during installation.

Procedure

1. Launch the installation wizard by double-clicking the setup.exe file in the \Windows subdirectory of the installation image.
2. Select **Next** on the Welcome window.
3. Click **Next** on the Select Features window without making any changes.
4. On the Agent Deployment window, select the agents that you want to add to the depot and click **Next**.
5. Review the installation summary and click **Next** to begin the installation. After the agents are added to the agent depot, a configuration window (called the Setup Type window) is displayed.
6. Clear all selected components. You have already configured all components on this computer and do not need to reconfigure any now. Click **Next**.
7. Click **Finish** to complete the installation.
8. Click **Finish** on the Maintenance Complete window.

Linux and UNIX: Populating the agent depot during installation

Follow steps in this section to populate the agent depot from the Linux or UNIX installation image.

Procedure

1. In the directory where you extracted the installation files, run the following command:
`./install.sh`
 2. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory (/opt/IBM/ITM). If you want to use a different installation directory, type the full path to that directory and press Enter.
 3. If the directory you specified does not exist, you are asked whether to create it. Type y to create this directory.
 4. The following prompt is displayed:
Select one of the following:
1) Install products to the local host.
2) Install products to depot for remote deployment (requires TEMS).
3) Install TEMS support for remote seeding
4) Exit install.
Please enter a valid number:

Type 2 to start the installation and press Enter.
The end user license agreement is displayed. Press Enter to read through the agreement.
 5. Type 1 to accept the agreement and press Enter.
 6. Type the number that corresponds to the agent or agents that you want to add to the agent depot and press Enter. If you are going to add more than one agent, use a comma (,) to separate the numbers.
To select all available agents, type all.
You can select multiple agents with consecutive corresponding numbers by typing the first and last numbers for the agents, separated by a hyphen (-). For example, to add all of the agents between 8 and 12, type 8-12.
To clear an agent that you previously selected, type the number for the agent again.
- Note:** Use the following keys to navigate the list of agents:
- | | |
|----------|-------------------------------------|
| U | Moves up a line in the list. |
| D | Moves down a line in the list. |
| F | Moves forward one page in the list. |
| B | Moves back one page in the list. |
7. When you have specified all the agents that you want to add to the agent depot, type E and press Enter to exit.

Populating the agent depot with the `tacmd addBundles` command

This section contains the `tacmd addBundles` command to populate the agent depot.

To populate the agent depot using the `tacmd addBundles` command, run the following command:

```
tacmd addBundles [-i IMAGE_PATH]
                  [-t PRODUCT_CODE]
                  [-p OPERATING_SYSTEM]
                  [-v VERSION]
                  [-n]
                  [-f]
```

For the full syntax, including parameter descriptions, see the *IBM Tivoli Monitoring Command Reference*.

Examples:

- The following example copies every agent bundle, including its prerequisites, into the agent depot on a UNIX computer from the installation media (CD image) located at /mnt/cdrom/:

```
tacmd addbundles -i /mnt/cdrom/unix
```

- The following example copies all agent bundles for the Oracle agent into the agent depot on a UNIX computer from the installation media (CD image) located at /mnt/cdrom/:

```
tacmd addbundles -i /mnt/cdrom/unix -t or
```

- The following example copies all agent bundles for the Oracle agent into the agent depot on a Windows computer from the installation media (CD image) located at D:\WINDOWS\Deploy:

```
tacmd addbundles -i D:\WINDOWS\Deploy -t or
```

- The following example copies the agent bundle for the Oracle agent that runs on the AIX version 5.1.3 operating system into the agent depot on a UNIX computer from the installation media (CD image) located at /mnt/cdrom/:

```
tacmd addbundles -i /mnt/cdrom/unix -t or -p aix513
```

Changing the location of the agent depot

By default, the agent depot is located in the *itm_installdir*/CMS/depot directory on Windows and *itm_installdir*/tables/*tems_name*/depot directory on UNIX. The **tacmd addBundles** command puts the agent bundle in that location unless another location is defined in the monitoring server configuration file for DEPOTHOME.

About this task

If you want to change this location, do the following before you run the **tacmd addBundles** command:

Procedure

1. Open the KBBENV monitoring server configuration file located in the *itm_installdir*\CMS directory on Windows and the *itm_installdir*/tables/*tems_name* directory on Linux and UNIX.
2. Locate the DEPOTHOME variable. If it does not exist, add it to the file.
3. Type the path to the directory that you want to use for the agent depot.
4. Save and close the file.
5. On UNIX or Linux only, add the same variable and location to the kbbenv.ini file located in *itm_installdir*/config/kbbenv.ini.

If you do not add the variable to the kbbenv.ini file, it will be deleted from the KBBENV file the next time the monitoring server is reconfigured.

Managing your agent depot

This section contains the commands for managing your agent depot.

Use the following commands to manage your agent depot:

Table 63. Agent depot management commands

Command	Description
tacmd listbundles	Lists the details for one or more bundles available to be added to the local agent depot.
tacmd removebundles	Deletes one or more bundles from the local agent depot.
tacmd viewdepot	Lists the types of bundles available in either the local or remote agent depot.

See the *IBM Tivoli Monitoring Command Reference* for the full syntax of these commands.

Note: Only Tivoli-provided product agent bundles should be loaded into the IBM Tivoli Monitoring deployment depot. User-provided or customized bundles are not supported. Use only Tivoli provided

tacmd commands to process bundles and to execute agent deployments. Manual manipulation of the depot directory structure or the bundles and files within it is not supported and may void your warranty.

Sharing an agent depot across your environment

If your monitoring environment includes multiple monitoring servers (a hub monitoring server and remote monitoring servers), you can put your agent depot in a central location, such as a shared file system, and access the depot from all of the monitoring servers.

After populating your agent depot with either of the methods described in “Populating your agent depot” on page 371, use the following steps to share the agent depot:

1. Open the KBBENV monitoring server configuration file located in the *itm_installdir*\CMS directory on Windows and the *itm_installdir*/tables/*tems_name* directory on Linux and UNIX.
2. Locate the DEPOTHOME variable. By default, the agent depot is located in the *itm_installdir*/CMS/depot directory on Windows and *itm_installdir*/tables/*tems_name*/depot directory on UNIX.
3. Type the path to the shared agent depot for the DEPOTHOME variable.
4. Save and close the file.
5. On UNIX or Linux only, add the same variable and location to the *kbbenv.ini* file located in *itm_installdir*/config/*kbbenv.ini*.

If you do not add the variable to the *kbbenv.ini* file, it will be deleted from the KBBENV file the next time the monitoring server is reconfigured.

If you are using a Windows monitoring server connecting to a depot on another Windows computer, you must set the service ID for the Windows monitoring server to "Administrator." Also, instead of specifying a mapped drive letter for the path to the depot directory, use the UNC path (such as \\server\share).

Use the following steps to change the service ID:

1. From the Control Panel, double-click **Administrative Tools**.
2. Double-click **Services**.
3. Right-click **Tivoli Enterprise Monitoring Svcs** and click **Properties**.
4. On the **Log On** tab, select **This Account**.
5. Type Administrator in the **This Account** field.
6. Type the password for the administrator in the **Password** field. Confirm the password by typing it again in the **Confirm password** field.
7. Click **Enable**.

If the Administrator user does not have Logon as a service right, you are prompted to add it.

Deploying OS agents

Before you can deploy any non-OS agent, you must first install an OS agent on the computer where you want the non-OS agent to be deployed. In addition to monitoring base OS performance, the OS agent also installs the required infrastructure for remote deployment and maintenance.

Note:

1. On Windows: IBM Tivoli Monitoring does not support more than one OS agent installation on the same machine. You can not use different directories to install more than one OS agent on the same Windows machine.
2. Ensure that you have populated your agent depot, as described in “Populating your agent depot” on page 371, before attempting to deploy any agents.

You can install the OS agent locally, as described in “Installing monitoring agents” on page 261 or remotely using the **tacmd createNode** command.

The **tacmd createNode** command creates a directory on the target computer called the *node*. The OS and non-OS agents are deployed in this directory. Agent application support is also added at this time (see “Installing and enabling application support” on page 275).

The **tacmd createNode** command uses one of the following protocols to connect to the computers on which you want to install the OS agent:

- Server Message Block (SMB), used primarily for Windows servers
- Secure Shell (SSH), used primarily by UNIX servers, but also available on Windows

Note: Only SSH version 2 is supported.

- Remote Execution (REXEC), used primarily by UNIX servers, but not very secure
- Remote Shell (RSH), used primarily by UNIX servers, but not very secure

You can specify a protocol to use; if you do not, the **tacmd createNode** command selects the appropriate protocol dynamically.

Requirements for the tacmd createNode command

This section contains the requirements for the **tacmd createNode** command.

Before you can use the **tacmd createNode** command to deploy OS agents, ensure the following requirements are met:

- The **createNode** command no longer has to be executed locally to the Tivoli Enterprise Monitoring Server.
- On both Windows, Linux, and UNIX, you must issue a **tacmd login** command prior to executing the **tacmd createNode** command.
- **Windows** The user ID that you specify using the **-u** parameter *must* have administrator privileges on the target computer.
- **Linux** **UNIX** You must specify the “root” user ID using the **-u** parameter and the root password using the **-p** parameter for the **tacmd createNode** command to execute correctly. No other user ID may be specified.
- Any computer to which you want to deploy the OS agent must have a supported protocol installed.
- Security in your environment must be configured to permit **createNode** to pass through the firewall, using the protocol that you specify in the command parameters.
- On Windows computers:
 - SMB requires that the default, hidden, and administrative shares be available on the drive being accessed and on the drive that hosts the System temporary directory.
 - SMB signing is not supported when connecting using SMB. The computer to which you are deploying an OS agent cannot require SMB signing.
 - For Windows XP, disable Simple File Sharing. Simple File Sharing requires that all users authenticate with guest privileges, which **createNode** does not support. To disable Simple File Sharing, perform the following steps:
 1. Open the Windows Explorer.
 2. Click **Tools → Folder Options**.
 3. Click the **View** tab.
 4. Scroll through the list of settings to **Use Simple File Sharing**.
 5. Clear the check box next to **Use Simple File Sharing** and click **OK**.
 - For Windows XP computers with Service Pack 2, disable the Internet Connection Firewall.
 - For Windows XP computers, set Network Access Sharing and Security to **Classic - local users authenticate as themselves**. Use the following steps:
 1. From the Control Panel, double-click **Administrative Tools**.

2. Double-click **Local Security Policy**.
 3. Expand **Local Policies** and click **Security Options**.
 4. Right-click **Network access: Sharing and security for local accounts** and click **Properties**.
 5. Select **Classic - local users authenticate as themselves** from the list and click **OK**.
- For all Windows computers, enable remote registry administration. (This is enabled by default.)
 - On UNIX systems, if you are using the RSH protocol, run the **tacmd createNode** command as root on the monitoring server.
 - If you are deploying the OS agent to a UNIX or Linux computer, that computer must have the ksh shell. Only the Korn shell is supported for the execution of the installation and runtime scripts.
 - If you are using SSH V2 (for either Windows or UNIX), configure SSH on the target computers to permit the use of password authentication. To permit this, do the following:
 1. Edit the `/etc/ss/sshd_config` file on the target computer.
 2. Locate the following line:
`PasswordAuthentication no`
 3. Change the no to yes and save the file.
 4. Restart the daemon.

Note: If you are using private key authentication in your environment, you do not need to set SSH to permit password authentication.

For more information see “Remote Execution and Access” on page 115.

Using the **tacmd createNode** command

To deploy an OS agent from the command-line interface, use the **tacmd createNode** command.

For example, the following command deploys the UNIX OS monitoring agent on the server1.ibm.com computer in the `/opt/IBM/ITM` directory. The installation is done as the root user. The option property `EXECPREREQCHECK=Y` performs a prerequisite check on the agent. For more information, see “Prerequisite Scanner for IBM Tivoli Monitoring agents” on page 48.

```
tacmd createNode -h server1.ibm.com -d /opt/IBM/ITM -u root -o EXECPREREQCHECK=Y
```

Important: Unless you specifically indicate otherwise, the agent that you deploy using this command assumes that the monitoring server to which it connects is the monitoring server from which you run the command. The agent also uses the default settings for the communications protocol (IP.PIPE for protocol type and 1918 for the port). To change these default values (especially if you are not using the IP.PIPE protocol), use the following property (specified with the `-p` parameter) when running the command: `SERVER=[PROTOCOL://][HOST|IP][:PORT]`. For example, `SERVER=IP.PIPE://server1.ibm.com:1918`.

The **tacmd createNode** command contains the optional `-k` group parameter. Passing the `-k` parameter when you remotely deploy an OS agent will execute the `secureMain` utility at the end of the deployment process to secure your IBM Tivoli Monitoring environment. For more information, see Appendix G, “Securing your IBM Tivoli Monitoring installation on Linux or UNIX,” on page 925. For the full syntax, including parameter descriptions, see the *IBM Tivoli Monitoring Command Reference*.

Deploying non-OS agents

Be aware of certain criteria for deploying non-OS agents.

You can deploy non-OS agents through the Tivoli Enterprise Portal or from the command-line.

Note:

1. The deployment and configuration of agents varies depending on the specific agent. The following procedures provide generic deployment information. For the exact values required for your agent, see the configuration information in the user's guide for the agent.
2. Ensure that you have populated your agent depot, as described in "Populating your agent depot" on page 371, before attempting to deploy any agents.
3. You must have already installed or deployed an OS agent on the computer where you are now deploying the non-OS agent and the agent must be running.
4. You can only deploy the Tivoli Performance Analyzer through the installation media, and not through the Tivoli Enterprise Portal or the command-line.

Deploying through the portal

Follow the steps in this section to deploy an agent through the portal GUI.

Before you begin

Before you deploy an agent through the Tivoli Enterprise Portal, application support for that agent must be installed on the portal server (see "Installing and enabling application support" on page 275).

Procedure

1. Open the Tivoli Enterprise Portal.
2. In the Navigation tree, navigate to the computer where you want to deploy the agent.
3. Right-click the computer and click **Add Managed System**.
4. Select the agent that you want to deploy and click **OK**.
5. Complete the configuration fields required for the agent. For information about these fields, see the configuration documentation for the agent that you are deploying.
6. Click **Finish**.
7. If the computer where you are deploying the agent already has a version of that agent installed, you can stop the deployment, add a new instance of the agent, if possible, or reconfigure the existing agent.

Results

A message will tell you when the deployment finishes successfully.

Deploying through the command-line

To deploy non-OS agents from the command-line, use the **tacmd addSystem** command.

See the *IBM Tivoli Monitoring Command Reference* for the full syntax of this command, including parameter descriptions. You can run the **cinfo** command (UNIX) or the **kincinfo -i** command (Windows) to list the product codes for agents installed on the current computer.

Each agent bundle has its own unique configuration parameters that you may need to specify using this command. If you have installed the agent bundle that you want to deploy to the deployment depot, you can view the configuration parameters by running the following command from the monitoring server where that agent bundle is installed:

```
tacmd describeSystemType -t pc -p platform
```

An agent of the same type and platform must be deployed into the depot available to the monitoring server from which the command is run. You can also get more information about agent-specific parameters in the agent user's guide for the agent that you want to deploy.

Note: The `tacmd` command has been updated to support two new features, the asynchronous remote deployment (which allows you to request the deployment of another remote agent, even if the previously deployed agent has not initialized fully) and the grouping of agents to be remotely deployed. The deployment commands have been modified to accept two new grouping parameters: the `-g` parameter lets you specify a deployment group and the `-b` parameter lets you specify a bundle group. For detailed information, including usage examples, see the *IBM Tivoli Monitoring Command Reference*.

With asynchronous remote deployment, CLI requests are queued, and the `tacmd` command returns control immediately to the process that invoked them along with a transaction ID that can be used to monitor the status of the request via either the `tacmd getDeployStatus` command or the Deployment Status Summary By Transaction workspace in the Tivoli Enterprise Portal. Asynchronous remote agent deployment applies both to agents started via the `tacmd` command and to those started using the portal client. Workspace reports regarding asynchronous agent deployment are also available; for information, see the IBM Tivoli Monitoring online help.

The prerequisite scanning feature allows you to check for agent installation prerequisites before performing an installation. The two mechanisms available are a manually executed stand-alone prerequisite scanner, or a remote prerequisite scanner facility that extends the capabilities of IBM Tivoli Monitoring's remote deployment component. For more information, see "Prerequisite Scanner for IBM Tivoli Monitoring agents" on page 48.

Deploying Netcool/OMNibus System Service Monitor (SSM) agents

Deploying Netcool/OMNibus SSM agents is very similar to deploying IBM Tivoli Monitoring OS agents: You first need to populate your agent depot with SSM bundles on the *Netcool/SSM/ASM Agents* DVDs or DVD images. There is DVD for each of the operating systems supported for these agents.

Once this is done, you can use the `tacmd createNode` command to deploy the SSM agent; see the *IBM Tivoli Monitoring Command Reference* for complete information on this command.

To populate the SSM bundles, download the SSM image from the Passport Advantage Web site, <http://www-01.ibm.com/software/howtobuy/passportadvantage/>, and use the `tacmd addBundles` command:

```
tacmd addBundles -t ssm -p li6213 -i c:\SSM40_Bundle
```

This example adds the SSM bundle for the Linux platform to the depot.

Once you have populated the depot with SSM bundles, you can use the `tacmd viewDepot` command to see the bundles. Each SSM bundle consists of the installer files and the corresponding descriptor file.

The `tacmd` commands for installing, patching, starting, stopping, and configuring the Tivoli Monitoring agents have been expanded to support operations with Netcool[®] SSM agents as well.

Installing an SSM agent

Use the `tacmd createNode` command to install an SSM agent.

The mechanism for pushing the SSM image to the agent machine is the same as that used to deploy a Tivoli Monitoring OS agent. Example:

```
tacmd createNode -h smb://achan1.raleigh.ibm.com -t SSM -u achan -d c:\SSMAgent\ssm -p SNMPPORT=16002
```

This example deploys an SSM agent to the Windows machine `achan1` using the Server Message Block. The installation directory for the agent is specified by the `-d` option, and the SNMP port to be used by the SSM agent is specified using the `-p SNMPPORT` property.

Uninstalling an SSM agent

To uninstall an SSM agent, use the **tacmd removeSystem** command.

Example:

```
tacmd removeSystem -h smb://achan1.raleigh.ibm.com -t SSM -u achan-d c:\SSMAgent\ssm
```

The above command uninstalls the SSM agent on Windows machine achan1.

Installing an SSM patch

Use the **tacmd updateAgent** command to patch an SSM agent.

1. Just like installing an SSM image, first run the **tacmd addBundles** command to add the patch images to the IBM Tivoli Monitoring depot.
2. Once this is done, run the **tacmd updateAgent** command to install the patch on the agent. The transfer of the patch image to the agent is done through the Tivoli Enterprise Monitoring Server's HTTP server.

Here is an example of an HTTP request for agent `http://TEMS_SERVER:1920//ms/kdh/lib/depot/ PACKAGES/WINNT\ssm\040000001\ssm40-fix pack1-win32-x86.exe`:

```
tacmd updateAgent -h achan1.raleigh.ibm.com -t SSM -l "fix pack 1"
```

This installs SSM 4.0 agent fix pack 1 to the achan1 agent machine. The agent must be up and running; if not, the request fails.

To start the agent, use the **tacmd startAgent** command, as explained in “Starting an SSM agent.”

Uninstalling an SSM patch

To uninstall an SSM patch, you also use the **tacmd removeSystem** command, with the `-l` option to remove the patch.

Example:

```
tacmd removeSystem -h achan1.raleigh.ibm.com -t SSM -l "fix pack 1" -p SNMP_PORT=16002
```

Starting an SSM agent

To start an SSM agent from IBM Tivoli Monitoring, use the **tacmd startAgent** command.

Example:

```
tacmd startAgent -h achan1.raleigh.ibm.com -p SNMP_PORT=16002
```

Stopping an SSM agent

To stop an SSM agent from IBM Tivoli Monitoring, use the **tacmd stopAgent** command.

Example:

```
tacmd stopAgent -h achan1.raleigh.ibm.com -p SNMP_PORT=16002
```

Restarting an SSM agent

To restart an SSM agent from IBM Tivoli Monitoring, use the **tacmd restartAgent** command.

Example:

```
tacmd restartAgent -h achan1.raleigh.ibm.com -p SNMP_PORT=16002
```

Configuring an SSM agent

There are several configuration functions you can perform with SSM agents using the **tacmd** **configureSystem** command.

- Change the logging file and level on the agent and perform a silent reboot of the agent. Example:

```
tacmd configureSystem -h achan1.raleigh.ibm.com  
-p SNMPPORT=16002 LogFile=test.log LogLevel=debug -r
```




- Transfer files, including executable files, to the agent machine. Example:

```
tacmd configureSystem -h achan1.raleigh.ibm.com -p SNMPPORT=16002 -l test1.exe test2.exe
```

- Transfer configuration files to the agent machine. Example:

```
tacmd configureSystem -h achan1.raleigh.ibm.com -p SNMPPORT=16002 -c test1.cfg test2.cfg
```

Note:

1. The files specified in the configfile (-c) option or the filelist (-l) option must be stored in the ssmconfig subdirectory of the depot directory.
 -  If your depot directory is c:\IBM\ITM\CMS\depot, you must put the files specified in the configfile option or the filelist option into the c:\IBM\ITM\CMS\depot\ssmconfig directory.
 -   If your depot directory is /opt/IBM/ITM/tables/TEMS_NAME/depot, put the files into the /opt/IBM/ITM/tables/TEMS_NAME/depot/ssmconfig directory.
2. The files are pulled from the agent through the Tivoli Enterprise Monitoring Server's HTTP server. Here is an example of an HTTP request:

```
http://tems_host_name:1920//ms/kdh/lib/depot/ssmconfig/test1.cfg
```

Bulk deployment of NetCool SSM agents

The bulk deployment of SSM agents is very similar to the bulk deployment of IBM Tivoli Monitoring agents.

See “Bulk agent deployment” on page 382 for more information.

Query deployment status of Netcool SSM agents

The same tools for checking the status of your IBM Tivoli Monitoring agent deployments also apply to SSM agents.

See “Deploy status” on page 384 for more information. Figure 82 on page 382 shows an example using the Tivoli Enterprise Portal's Deployment Status Summary workspace.

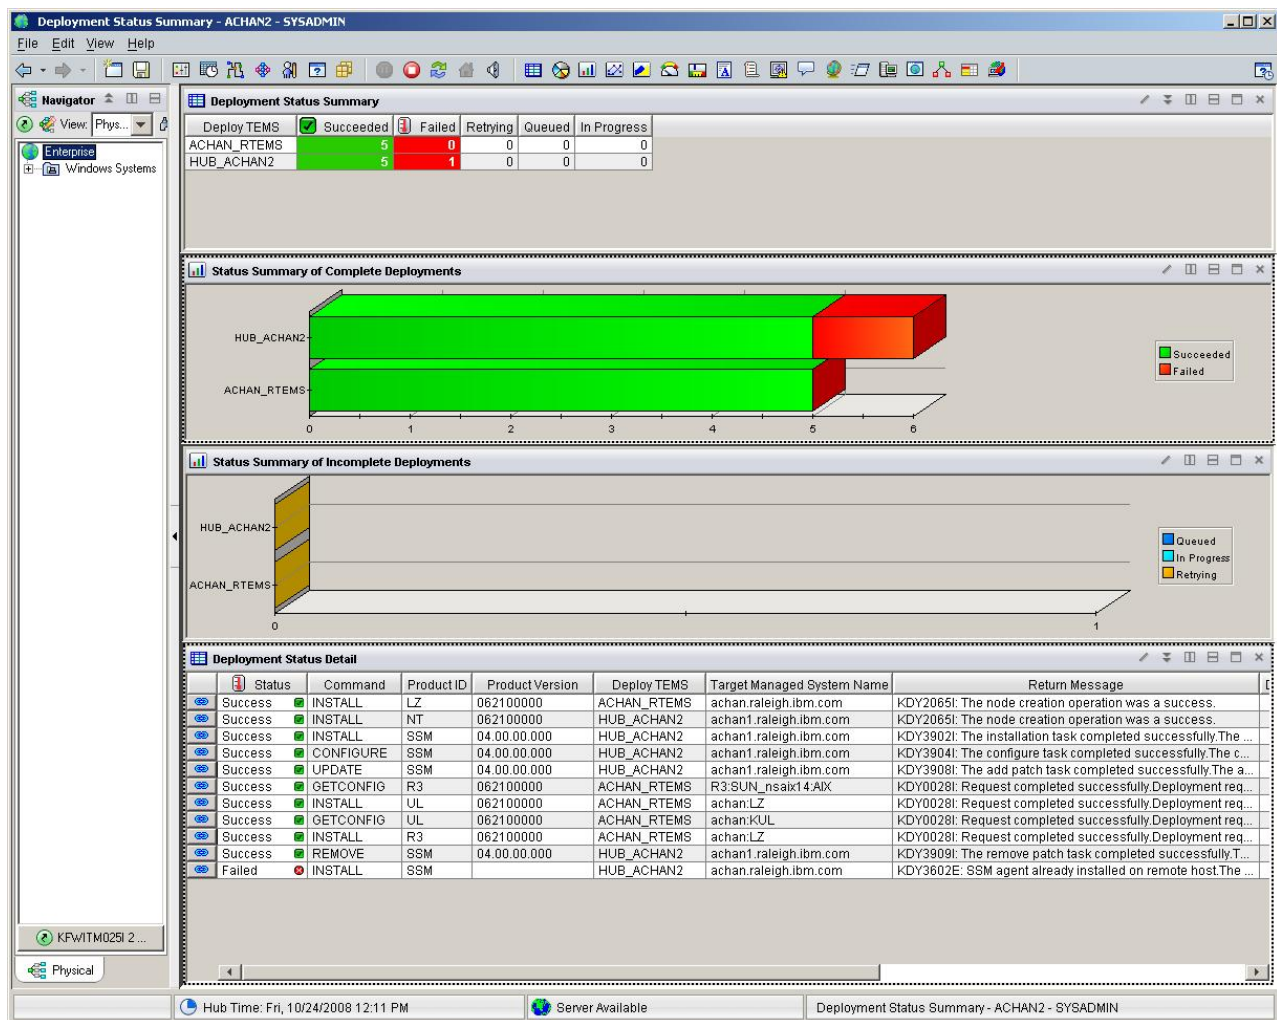


Figure 82. Deployment Status Summary workspace showing the status of SSM deployments

Bulk agent deployment

Bulk deployment of IBM Tivoli Monitoring agents into moderate to large environments is made possible by a set of capabilities and best practices coming together to achieve rapid rollout with ease.

The capabilities that enable bulk deployment in larger environments fall into two main categories:

- The components involved in deployment.
- The processing model associated with deployment.

The best practices described in this section are recommended or common ways that you can use these capabilities to successfully deploy agents in larger environments.

Deployment processing model

The processing model used during deployment is critical to the success and timeliness of deployments.

The two critical aspects of deployment processing are:

- Asynchronous transaction submission on the front end.
- Parallel transaction handling on the back end.

The asynchronous nature of the IBM Tivoli Monitoring deployment model allows you to submit deployment transactions into the environment quickly without having to wait for the deployment request to complete. Instead of the deployment command returning after some period of time with a message saying that the deployment completed with a success or a failure, the CLI returns immediately with a transaction identifier (ID). This frees your user environment (be it the Tivoli Enterprise Portal, the command-line, or an automation script) to continue normal processing:

```
C:\>tacmd createnode -g Targets -b Agents
```

```
KUICCN022I: Request has been successfully queued to the deploy controller.  
The transaction ID is 1224009912765000000000041, use the getDeployStatus  
CLI to view the status.
```

You can then monitor the status of your deployment requests using the command interface or Tivoli Enterprise Portal workspaces or situations to check for failures.

Once a set of deployment transactions has been submitted into the IBM Tivoli Monitoring infrastructure, they are routed to the appropriate remote Tivoli Enterprise Monitoring Server for processing. Each remote monitoring server can process deployment transactions independently of every other monitoring server, whether remote or hub. This allows for a large degree of parallelism, which increases with each remote monitoring server added to your environment. Additionally, each remote Tivoli Enterprise Monitoring Server can process multiple deployment transactions concurrently, further increasing the level of parallelism applied to bulk deployments.

The default number of concurrent deployments per monitoring server is 10, but this is configurable using the `DEPLOYTHREADPOOLSIZE` setting in the Tivoli Enterprise Monitoring Server configuration file (`KBBENV`).

Figure 83 illustrates the processing model for bulk deployment.

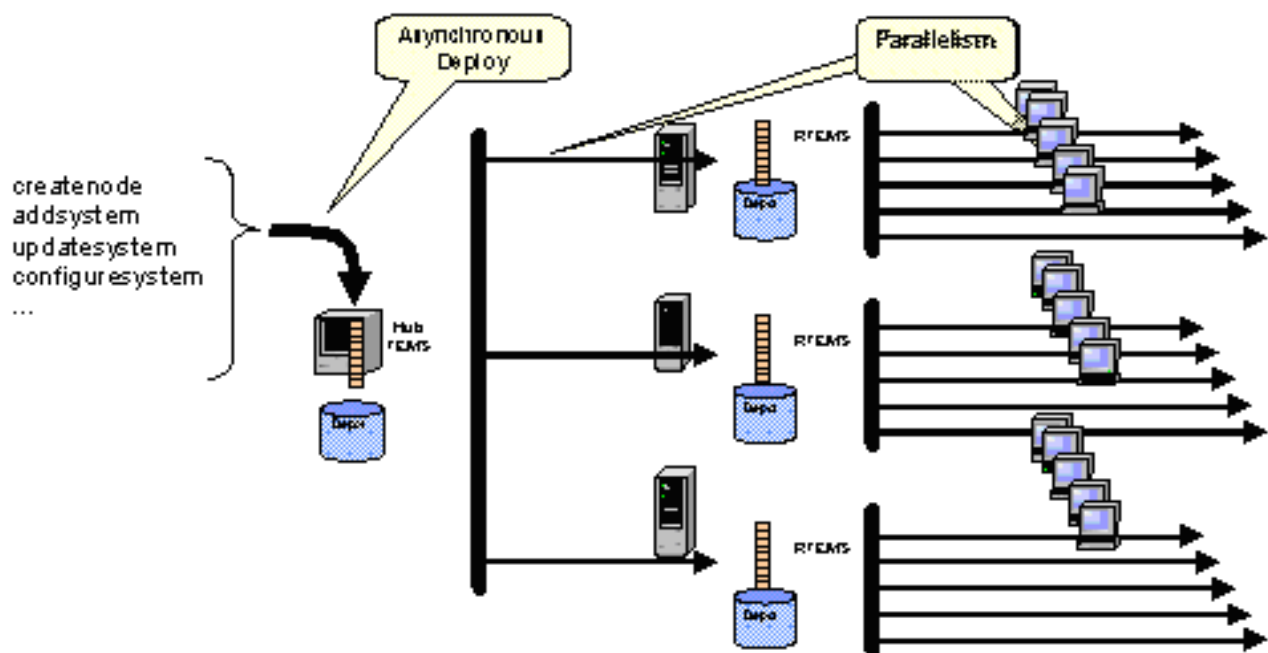


Figure 83. Bulk deployment processing model

Deploy status

Because deployment transactions are asynchronous (that is, they run in the background), monitoring the status of deployments is an important activity to ensure they complete successfully.

There are three mechanisms for monitoring deployment status, as follows.

- The **tacmd getDeployStatus** CLI command.
For more information about this command, see the *IBM Tivoli Monitoring Command Reference*.
- The Tivoli Enterprise Portal's Deployment Status workspaces.
- Deployment attribute groups, and situations created using the Tivoli Enterprise Portal that are based on these groups.

Using these features, you can display and monitor the progress of deployments running in your environment. You can track the status of your deployments either by individual transaction or by group transaction. A group transaction represents a deployment using groups of agents deployed to groups of targets; this is discussed in the succeeding sections. Each transaction possesses a status that indicates the progress the transaction has made in the deployment process. The valid status values for a deployment transaction are listed below.

Status Description

Pending

The transaction is waiting in the queue to begin processing.

In-Progress

The transaction has begun the deployment process but has not yet finished.

Success

The transaction has completed successfully.

Failed Retrying

The transaction experienced a recoverable error during the deployment process and will be restarted periodically at the point of failure for a predefined number of iterations. The number of retry iterations is defined by configurable property `ERRORRETRYLIMIT`. The time between each successive attempt is defined by the `RETRYTIMEINTERVAL` property. Both of these properties can be set in the monitoring server's configuration file, `KBBENV`, and take effect when the monitoring server is restarted. The default value for `ERRORRETRYLIMIT` is 3 retries, and the default value for `RETRYTIMEINTERVAL` is 7 minutes. These default values are based on a typical bulk deployment scenario using the `createNode` command. A smaller `RETRYTIMEINTERVAL` value (for example, 20 seconds) might be better suited if you perform more bulk upgrade scenarios than deployments.

If an error persists after all retry iterations are exhausted, the transaction moves to Failed status.

Failed The transaction has completed with an unrecoverable failure, or the number of retry attempts has been exhausted. Consult the status message included with the transaction status for more information about the failure.

Deployment Status workspaces

There are six workspaces related to bulk deployment. Three can be accessed directly from the workspaces menu of the Navigator tree's Enterprise item. The remaining three can be accessed by selecting links.

These workspaces are:

- Deploy Depot Package List
- Deployment Status Summary
- Deployment Status Summary by Transaction ID
- Deployment Status Summary by Product
- Deployment Status Summary by Deploy Group
- Installation Logs

For more information about these workspaces, see the *IBM Tivoli Monitoring Tivoli Enterprise Portal User's Guide* or the Tivoli Enterprise Portal online help.

Deployment attribute groups and situations

The hub Tivoli Enterprise Monitoring Server provides two attribute groups that you can use to monitor the status of bulk deployments within your IBM Tivoli Monitoring environment: Deploy Status and Deploy Summary.

You can use these attribute groups to create situations that notify you in the event of a failure or exceptional circumstance related to a deployment. For more information about these deployment status attribute groups, see the *IBM Tivoli Monitoring Tivoli Enterprise Portal User's Guide* or the Tivoli Enterprise Portal online help.

IBM Tivoli Monitoring provides two deployment situations that monitor the Deploy Status attribute group, which you can use directly or as samples for creating your own situations. These situations are as follows.

Situation

Description

Deploy_Failed

Triggers a critical event when a deployment transaction fails.

Deploy_Retrying

Triggers a minor event when a deployment transaction enters the failed_retry state.

Organizing deployments using groups

There are three primary components that come into play when performing bulk deployments. These are 1) the components being deployed, 2) the recipients (*targets*) of the deployment, and 3) properties that specify, clarify, or alter key details used during the deployment.

Planning and organizing these components in a manner that reflects the current condition of your environment as well as the desired outcome for your environment will result in successful bulk deployments.

Organization of deployment targets and deployable components is achieved using *groups*. There are two types of groups used in bulk deployment, as follows.

Type Function

Deploy groups

Organize deployment targets, which are the systems or agents to which a deployment is targeted.

Bundle groups

Organize the deployable components that will be used during deployment operations.

In addition to the obvious role these groups play in organizing the participants of deployment, the grouping facilities perform another very important role: Groups and group members can also hold and organize properties that the deployment process uses; this is discussed in the following sections.

Deploy and Bundle groups are created using `tacmd` commands. For detailed information about creating such groups, see the *IBM Tivoli Monitoring Command Reference*.

Deploy groups

Deploy groups organize the targets of deployment operations. Deploy targets are specified as members within the deploy group. These deploy targets can represent unmanaged systems, which could be targets of a `createNode` deployment operation. Deploy targets can also represent managed systems where a Tivoli Monitoring or Netcool System Service Monitor (SSM) agent is already installed. Both of these types of deploy targets can be added as members to deploy groups.

This dual representation of targets within deploy groups allows a group to specify:

- Targets of an initial OS agent or SSM deployment using the **tacmd createNode** command.
- Targets of an application agent deployment using the **tacmd addSystem** command.
- Targets of an agent or SSM upgrade using the **tacmd updateAgent** command.
- Targets of an agent or SSM configuration deployment using the **tacmd configureSystem** command.

Once you create the necessary deploy groups to represent logical groupings of systems within your enterprise, these groups can be used repeatedly during the lifecycle of your monitoring infrastructure to deploy, update, start, stop, configure, and eventually remove agents.

Deploy groups best practices:

There are a number of best practices associated with deploy groups.

The following are suggested:

- Segment your systems by geographic location. In this case, it is likely that you would want all agents deployed to a deploy group to connect with the same primary Tivoli Enterprise Monitoring Server.
- Segment your systems according to their role within the enterprise. You might create a group of database servers, a group of application servers, a group of Web servers. This allows you to deploy and administer agents to the systems within a group the same way but differently from how you might deploy and administer agents in a different group.
- Segment your systems according to the monitoring infrastructure being used. You might create a set of groups for monitoring using Tivoli Monitoring agents and another group or set of groups for monitoring using SSM agents.
- Segment your systems into smaller groups representing a percentage of total systems. This allows you to perform deployments into more easily consumable sets of activities.

Bundle groups

Bundle groups organize and represent the components that you deploy. These include IBM Tivoli Monitoring agents and Netcool SSM agents. Deployable bundles are added as members to bundle groups using the **tacmd** command by specifying the agent's product code and optionally the platform and version for the bundle.

To enable bulk deployment of more than one agent at a time, you add multiple members to the bundle group representing multiple agents to deploy. For example, you might create a bundle group to deploy a Windows OS agent and a DB2 application agent by adding a member with product code NT and a member with product code UD. Since one of these members is an OS agent, you use the **tacmd createNode** command to deploy this bundle group.

If you had created a bundle group containing only application agents, such as DB2, Domino, or SAP agents, you deploy them using the **tacmd addSystem** command. Each bundle group member represents a deployment bundle that must be added to the appropriate agent depot using the **tacmd addBundles** command before a bulk deployment command is executed involving the bundle group. To avoid the complexity of managing multiple depots, consider using a shared agent depot; see "Sharing an agent depot across your environment" on page 375.

Bundle groups best practices:

There are a number of best practices associated with bundle groups.

The following are suggested:

Grouping

Definition

Grouped by platform type

A group used to specify one or more agent bundles for a specific platform type.

Grouped by system role

A group used to specify a standard set of agents being deployed for a specific machine type or role within your enterprise. For example, a group deployed to database servers could contain an OS agent and a database agent.

Grouped for latest update

A group containing various agents where the version specification for agent member bundles is omitted. This causes initial deployment and upgrades to always deploy the most current available version in the depot. By managing the version of agents in the depot, this type of group always deploys the latest versions to target systems.

Group properties

When using the `tacmd` command for deployment (`createNode`, `addSystem`, etc.), you can use the `-p` option to specify properties that modify parameters of the deployment or to specify configuration parameters to the agent installation process. When using the command interface with `deploy` and `bundle` groups, properties previously specified on the `deploy` command-line can be attached to the groups themselves or to group members.

Attaching properties to groups and group members means that they are available for reuse each time the group is referenced in a deployment command without your needing to specify properties again on the command-line. This creates a very powerful mechanism for defining behaviors around your deployments and reusing them each time you specify a group as part of a deployment.

When creating your `deploy` and `bundle` groups for bulk deployment, you can assign properties to these groups. Properties at the group level apply to all members of the group. Additionally, properties can be assigned to individual group members. Any property assigned to a member applies to that member only and overrides a similar property specified at the group level. This allows you to build a hierarchy of properties that is consolidated during deployments to create a complete deployment request. Properties specified on the groups and group members can include any combination of OS agent, application agent, and deployment properties. See the *IBM Tivoli Monitoring Command Reference* for more information about specific properties.

Consider the following example where you want to deploy the UNIX OS agent to five AIX systems with the following hostnames.

```
aix1.mycompany.com
aix2.mycompany.com
aix3.mycompany.com
aix4.mycompany.com
aix5.mycompany.com
```

To do this, you must provide login credentials. All five systems share the same root userid and password. When creating the group, you can assign the userid and password to the group so it applies to all members. This way, you need specify the login credentials only once.

```
# KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=mypass

# tacmd addgroupmember -t deploy -g Targets -m aix1.mycompany.com
# tacmd addgroupmember -t deploy -g Targets -m aix2.mycompany.com
...
```

By attaching the userid and password to the `deploy` group at the group level, the properties apply to all members without having to repeat them for each member.

Suppose that one of the systems was on the far side of a slow link and required more time to complete a deployment. By adding a property to that member, you can change the behavior associated with that member without affecting the behavior of the others.

```
# tacmd addgroupmember -t deploy -g Targets -m aix3.mycompany.com -p
#KDYRemote Execution and Access.TIMEOUT=3600
```

But suppose that one of the members had a different root password. By attaching this password to the member for that system, you override the value at the group level and replace it with the member-specific value.

```
# tacmd addgroupmember -t deploy -g Targets -m aix5.mycompany.com -p
#KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=yourpass
```

Table 64 illustrates how the properties from the above example combine to build a useful set of consolidated properties.

Table 64. Interaction between member properties and group properties

Member	Group properties	Member properties	Consolidated properties
aix1	KDYRemote Execution and Access.Remote Execution and AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=myspass		KDYRemote Execution and Access.Remote Execution and AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=myspass
aix2	KDYRemote Execution and Access.Remote Execution and AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=myspass		KDYRemote Execution and Access.Remote Execution and AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=myspass
aix3	KDYRemote Execution and Access.Remote Execution and AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=myspass	KDYRemote Execution and Access.TIMEOUT=3600	KDYRemote Execution and Access.Remote Execution and AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=myspass KDYRemote Execution and Access.TIMEOUT=3600
aix4	KDYRemote Execution and Access.Remote Execution and AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=myspass		KDYRemote Execution and Access.Remote Execution and AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=myspass
aix5	KDYRemote Execution and Access.Remote Execution and AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=myspass	KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=yourpass	KDYRemote Execution and Access.Remote Execution and AccessUSERNAME=root KDYRemote Execution and Access.Remote Execution and AccessPASSWORD=yourpass

As you can see in the above example, group and member properties are joined during deployment to create a consolidate set of properties that are used for the deployment. Notice, however, that the deploy group member properties take precedence over and can actually override those at the deploy group level. This is visible with aix5.mycompany.com where the member value for KDYRemote Execution and Access.Remote Execution and AccessPASSWORD overrode the group value.

This same property mechanism applies to bundle groups as well. There might be a property that you want to apply to an agent bundle in a group regardless of where they are deployed. In this case you

would attach the property to the bundle group. Likewise, you might want to specify a property that applies to a single agent type within a bundle group. In this case you attach the property to the bundle member.

During a deployment operation where you deploy a bundle group to a deploy group, property consolidation occurs across all properties at the group and member level for both groups. Table 65 lays out the precedence of group/member properties.

Table 65. Property precedence between deploy groups and bundle groups

Precedence	Property location	What goes here
highest	Deploy group members	Properties specific to the individual member target.
	Deploy group	Properties common to all targets of the group. For example, if agents deployed to every target member of the group will connect to the same remote monitoring server, then you can specify the server's connection properties here.
	Bundle group member	Properties common to all deployments of the specific member bundle of the group. For example, if the bundle group contains the DB2 agent for Windows and you want all installations of DB2 agents on Windows to use a specific property, specify it here.
lowest	Bundle group	Properties common to all deployments of every bundle within the group, regardless of target.

Best-practice deployment procedures

The following sections provide some best-practice procedures to assist you in planning and implementing your bulk deployment.

For more information about the commands listed in the sections below, see the *IBM Tivoli Monitoring Command Reference*.

Deployment planning and preparation

Employ the best practices in this section in advance of actual deployments. They represent the setup steps necessary to prepare for actual deployment operations.

Procedure

1. Add bundles to the deploy depot.

```
# tacmd addbundles -I imagepath ...
```

2. Create deploy groups.

Create a deploy group for any collection of deployment targets. Group properties can be applied to the group during creation or afterward. Use any of the best practices described above in “Deploy groups” on page 385.

```
# tacmd creategroup -t DEPLOY -g DBServers ...
```

3. Add deploy group members.

After creating your deploy groups, add the appropriate systems hostnames as group members (targets) according the best practice chosen for the group. Member properties can be applied during member creation or afterward.

```
# tacmd addgroupmember -t DEPLOY -g DBServers -m hostname ...
# ...
```

4. Create bundle groups.

Create a bundle group for any collection of agent bundles you wish to deploy. Group properties can be applied to the group during creation or afterward. Use any of the best practices described above in “Bundle groups” on page 386.

```
# tacmd creategroup -t BUNDLE -g DBAgentBundles ...
```

5. Add bundle group members.

After creating your bundle groups, add the appropriate agent bundles as members to the groups according the best practice chosen for the group. Member properties can be applied during member creation or afterward.

```
# tacmd addgroupmember -t BUNDLE -g DBAgentBundles -m UnixOSAgent -y UX ...
# tacmd addgroupmember -t BUNDLE -g DBAgentBundles -m DB2Agent -y UD ...
# ...
```

Deployment

The following best practices are employed to initiate and perform actual deployments. These best practices apply regardless of whether the deployment is an initial agent installation, an agent upgrade, or an agent configuration update.

About this task

It is important to recognize that the groups created for deployment are applicable to all agent maintenance operations throughout the agent lifecycle.

Procedure

1. Submit deployment.

Submitting the deployment differs slightly depending on the type of deployment activity you desire, but they are all very similar.

- **createNode**

Use the **tacmd createNode** command when performing initial deployments of OS and SSM agents. This includes deploying bundle groups that contain OS agents grouped together with application agents.

```
# tacmd createnode -g DBServers -b OS_DBAgentBundles
```

- **addSystem**

Use the **tacmd addSystem** command when initially deploying application agents.

```
# tacmd addsystem -g AppServers -b AppAgentBundles
```

- **updateAgent**

Use the **tacmd updateAgent** command when performing agent or SSM upgrades.

```
# tacmd updateagent -g DBServers -b LatestAgentfix packs
# tacmd updateagent -g AppServers -b LatestAppAgentfix packs
```

Note: The KUIWINNT.dsc file on Windows systems and the *uiplatform.dsc* files (where *platform* is the platform name) on Linux and UNIX systems have been added or updated so that the KUI package (the tacmd commands) can be remotely deployed. Use the following command:

```
tacmd updateAgent -t ui -n node
```

where:

node

is the IBM Tivoli Monitoring node on which the KUI package is to be remotely deployed.

- **configureSystem**

Use the **tacmd configureSystem** command when performing configuration updates to agents or SSMs.

```
# tacmd configuresystem -g DBServers -b DBServerAgentConfig
```

- **removeSystem**

Use the **tacmd removeSystem** command when removing application agents or SSM patches from a system.


```
# tacmd removesystem -g AppServers -b AppAgentBundles
```

2. Monitor deployment.

There are many mechanisms that you can use for monitoring deployment status. These include:

- Using the CLI.

The **tacmd getDeployStatus** command returns the complete status of requested deployment transactions.

```
# tacmd getdeploystatus -g group_transactionID
```

- Using portal workspaces.

The Tivoli Enterprise Portal provides a number of workspaces that display the status of deployments. These provide an easy mechanism to rapidly visualize deployment status summaries as well as detailed status. See the or the portal online help for more information.

- Using automation.

Automation can be driven in two ways. Situations can be created to indicate deployment failures to initiate corrective activities. Also, automation scripts can use the **tacmd getDeployStatus** command to retrieve and parse status information and initiate corrective activities. Situations and automation scripts can be used in tandem to monitor deployment status.

- Regardless of the specific mechanism for monitoring deployment status, the best practice involves the following steps.

- a. Check the transaction status to ensure it is queued properly.
- b. Wait some period of time for the deployments to proceed.
- c. Check the transaction status for successful completions and failures. This involves either using the command interface or the portal interface or checking for situation events that indicate deployment failures.
- d. Perform any necessary corrective actions.
- e. Repeat starting with step 2b until all transactions have completed.

3. Correct problems.

When a deployment transaction fails, some corrective action is usually required before reinitiating the deployment. Consult the return message portion of the transaction status to identify the problem and possible resolutions. Rerun the deployment once errors are corrected. For example, when using a deploy group with the **addSystem** or **updateAgent** commands, remote deployment may fail to locate the existing managed-system name for some hosts. The message received is **KDY0012E** The target *target_host_name* is incorrect or is offline. The command did not complete because the value for the target is incorrect or the target is offline.

This message normally indicates that the OS agent is not online. If the agent is, in fact, online, cancel current operations to this node:

```
# tacmd cleardeploystatus -h hostname
```

Then issue the operation directly by using the managed system name parameter instead of the deploy group:

```
# tacmd updateAgent -t product_code -n managed_OS
```

4. Purge status.

When all deployment transactions have completed or the deployment status is no longer required, clear the transaction's deployment status from the systems.

```
# tacmd clearDeployStatus -g group_transactionID
```

Results

While the steps above describe the complete process of deployment, it is common to perform deployments in small sets of consumable transactions. That is, perform an initial deployment on a smaller group of systems to ensure that there are no issues that would inhibit the successful completion of deployment. Once that deployment is complete, continue with one or more larger deployments using

target deploy groups with more members. Work through each deploy group representing some portion of your enterprise until you have completed all deployments for your enterprise.

Working with non-agent bundles

This section describes how to deploy, update, and remove non-agent bundles.

Traditionally, remote deployment supported remote installation of deployable agents to a connected monitoring server. Now you can use remote deployment to install custom bundles or bundles for other components. This non-agent bundle capability is a way to deploy components that are not required to connect to the monitoring server. This enhancement includes the following:

- The allowable length for a product code for non-agent bundles is restricted to be between 3 to 32 characters in order to distinguish it from the agent product codes.
- The following **tacmd** commands were enhanced to support product codes up to 32 characters: **addBundles**, **addSystem**, **getDeployStatus**, **listBundles**, **removeSystem**, **updateAgent**, and **viewDepot**.

Note:

1. Avoid non-agent bundle product codes that start with K as this name might conflict with IBM Tivoli Monitoring agent bundles.
2. Non-agent bundles do not support the following **tacmd** commands: **startAgent**, **stopAgent**, **restartAgent**, **viewAgent**, and **configureSystem**.

Deploying a non-agent bundle

Complete the steps in this section to deploy a non-agent bundle.

Procedure

1. Run the **tacmd listBundles** command to view the agent bundles.
2. Add the agent bundle to the depot using the **tacmd addBundles** command.
3. Run the **tacmd viewDepot** command to display the contents of the agent depot. Verify the non-agent bundles were added to the depot.
4. Run the **tacmd addSystem** command with appropriate configuration information. If you want to specify an installation path, run this command with the **KDY.installPath** property set to **CANDLEHOME/productcode** if you want the location to be within the **CANDLEHOME** directory, or set to **/productcode** if you do not want the location to be within the **CANDLEHOME** directory, for example:

```
tacmd addSystem -p KDY.installPath=/IBM/ITM/Uxxx -n managed_os -t Uxxx
```
5. Check to ensure the status of transactions is queued or InProgress.

What to do next

Monitor the workspace or use the **tacmd getDeployStatus** command to verify the result. When the deployment completes successfully, verify that the bundle was installed at the default location if you did not specify the installation path, or to the location that you chose if you did specify an installation path.

Note:

1. The OS agent must already be installed on the system before deploying the non-agent bundle.
2. The non-agent bundle must exist in the depot of the monitoring server to which the OS agent is connected.
3. For more information about each of these commands, see the table below that links to a description of each command, including syntax and option information, along with examples.

Updating a non-agent bundle

Complete the steps in this section to update a non-agent bundle.

Procedure

1. Run the **tacmd listBundles** command to view the agent bundles.
2. Add the agent bundle to the depot using the **tacmd addBundles** command.
3. Run the **tacmd viewDepot** command to display the contents of the agent depot. Verify the non-agent bundles were added to the depot.
4. Run the **tacmd updateAgent** command with appropriate parameters. If the non-agent bundle was installed in a directory other than the default, run this command with the `KDY.installPath` property set to `CANDLEHOME/productcode` if the location is within the `CANDLEHOME` directory, or run this command with the `KDY.installPath` property set to `/productcode` if the location is not within the `CANDLEHOME` directory, for example:

```
tacmd updateAgent -p KDY.installPath=/IBM/ITM/Uxxx -n managed_os -t Uxxx
```
5. Check to ensure the status of transactions is queued or `InProgress`.

What to do next

Monitor the workspace or use the **tacmd getDeployStatus** to verify the result. When the deployment completes successfully, verify that the agent bundles were updated on the target system.

Note:

1. Ensure that the installation path is the same one that was used when the non-agent bundle was added initially.
2. The non-agent bundle must exist in the depot of the monitoring server to which the OS agent is connected.
3. For more information about each of these commands, see the table below that links to a description of each command, including syntax and option information, along with examples.

Removing a non-agent bundle

Follow the steps in this section to remove a non-agent bundle.

Procedure

1. Run the **tacmd removeSystem** command with appropriate parameters, including the `KDY.installPath` property if the agent was installed at a user specified location, for example:

```
tacmd removesystem -p KDY.installPath=/IBM/ITM/Uxxx -n managed_os -t Uxxx
```
2. Check to ensure the status of transactions is queued or **InProgress**.

What to do next

Monitor the workspace or use the **tacmd getDeployStatus** command to verify the result. When the removal completes, verify that the non-agent bundle files are removed from the target system and none of the IBM Tivoli Monitoring system files were inadvertently removed.

Note:

1. The non-agent bundle must already be in the depot in order to remove it. If it is not, add it to the depot using the **tacmd addBundles** command.
2. Ensure that the installation path is the same one that was used when the non-agent bundle was added initially.
3. The non-agent bundle must exist in the depot of the monitoring server to which the OS agent is connected.

Running deployment in a Hot Standby environment

The IBM Tivoli Monitoring Hot Standby capability allows your monitoring environment to continue operating in the event of environmental or operational issues with the primary hub Tivoli Enterprise Monitoring Server.

For detailed information about Tivoli Monitoring's Hot Standby feature, see the *IBM Tivoli Monitoring High Availability Guide for Distributed Systems*. You should refrain from deploying or updating agents when IBM Tivoli Monitoring is converting to a mirror monitoring server. No agent deployments or remote deployment operations should be executed from a Hot Standby mirror hub, as this may cause your deployment transactions to get stuck in a queued state, and you may not be able to clear them.

Self-describing agent installation

Product support files for IBM Tivoli Monitoring agent applications must be installed at each IBM Tivoli Monitoring server component. The self-describing agent feature combines the installation of an agent with the automatic dispersal and installation of the associated product support files throughout your IBM Tivoli Monitoring infrastructure.

The self-describing agent feature must be enabled on the following IBM Tivoli Monitoring components:

- Hub Tivoli Enterprise Monitoring Server.
- Tivoli Enterprise Portal Server (self-describing agent capability is enabled by default).
- Self-describing agent supported agents (self-describing agent capability is enabled by default).
- Any Remote Tivoli Enterprise Monitoring Server the agent is connected to (self-describing agent capability is enabled by default).

Enabling the self-describing agent capability at the hub monitoring server controls the capability across all components. For more information, see “Enabling self-describing agent capability at the hub monitoring server” on page 204.

Managing your self-describing agent installations

You can use `tacmd` commands to exercise more control over the self-describing agent products that are installed by the self-describing agent process.

As of IBM Tivoli Monitoring V6.3, the hub monitoring server does not allow the automatic installation of a self-describing agent enabled product. You must first specify the products and versions that can be installed, or change the default self-describing agent installation behavior to permit installations for all products and versions.

The following table describes the administration commands for managing your self-describing agent installations:

Table 66. tacmd commands to manage your self-describing agent installations

tacmd Command	Description
<code>tacmd addsdainstalloptions</code>	Specify the products and versions that can be installed by the self-describing agent facility.
<code>tacmd editsdainstalloptions</code>	Modify an existing product version configuration, or change the default self-describing agent installation behavior.
<code>tacmd listsdainstalloptions</code>	Display the existing self-describing agent installation configurations from the hub monitoring server.
<code>tacmd deletesdainstalloptions</code>	Remove a specific version configuration, or remove all versions for a product.

Beginning with V6.3, if you want to permit self-describing agent installations for all products and

versions without any blocking, you must perform the following steps:

1. Ensure that there are no products configured for self-describing agent installation by issuing the command `tacmd deletesdainstalloptions -all`.

Attention: This will remove any previous self-describing agent product versions and configurations that were added.

2. Issue the command `tacmd editsdainstalloptions -t DEFAULT -i ON`.

Self-describing agent installations for all products and versions without any blocking is essentially the default behavior in V6.2.3 and V6.2.3 FP1. To use the V6.3 default behavior and block automatic installations, you can issue the command `tacmd editsdainstalloptions -t DEFAULT -i OFF`. You can then permit the installation of specified products and versions by using the `tacmd addsdainstalloptions` command.

For the full syntax, including parameter descriptions, see the *IBM Tivoli Monitoring Command Reference*.

The following IBM Tivoli Monitoring components are not updated by self-describing agent installations:

- The Tivoli Enterprise Monitoring Automation Server component contains the Open Services Lifecycle Collaboration Performance Monitoring (OSLC-PM) service provider. The Performance Monitoring server provider does not support the dynamic refresh of application support files. This means that the Tivoli Enterprise Monitoring Automation Server must be recycled when new application support is added to a system, regardless of whether the support was added through a self-describing agent, a regular installation, or a manual copy. However, if the monitoring agent does not provide OSLC support, then the Automation Server does not need to be recycled. Operating system agents before IBM Tivoli Monitoring V6.3 do not provide OSLC support. For other types of monitoring agents, refer to the agent documentation to determine if they provide OSLC support.
- If the self-describing agent feature installs application support on the Tivoli Enterprise Portal Server and the dashboard data provider component is enabled, the portal server must be restarted so that the dashboard data provider can pick up the new or changed application support.
- The Tivoli Enterprise Portal desktop client

Related concepts:

“Dynamically controlling the hub monitoring server self-describing agent capability” on page 205
When the hub Tivoli Enterprise Monitoring Server is configured for self-describing agent capability (by setting the variable `KMS_SDA=Y`) and started, you can dynamically change the operational state of the hub self-describing agent feature without recycling the hub monitoring server.

Specifying what products and versions are installed on your monitoring server and portal server by the automatic self-describing agent process

Follow the steps in this section to specify what products and versions are installed on your monitoring server and portal server by the automatic self-describing agent process.

Procedure

1. Upgrade your Tivoli Monitoring environment to V6.3.
2. Determine what IBM Tivoli Monitoring products and versions you are ready to install on your environment.
3. Run the following command to define the products and versions the self-describing agent facility is allowed to install:

```
tacmd addsdainstalloptions -t xx -v 06230100
```

where:

`xx` Is the product code.

`06230100`

Is the version string.

4. Run the following command to view the current installation configurations:
`tacmd listsdainstalloptions`
5. If the hub monitoring server is configured for self-describing agent capability (by using the configuration variable `KMS_SDA=Y`), start the self-describing enabled agents.
6. Run the following command to view the status of the self-describing agent product installation:
`tacmd listappinstallrecs`

Terminal self-describing agent installation errors

In this section you can find examples of terminal self-describing agent installation errors.

In this example, the install record for product code 11 displays a STATE value of **ME**:

HUB/RTEMS	PRODUCT	VERSION	GRPID	ID	IDVER	SEEDSTATE	STATE	STATUS
RTEMS_LZ	11	06230000	5655	TMS	06230000	Y	ME	1014

The STATE=ME record indicates a self-describing agent metadata installation error has occurred on the monitoring server. The monitoring server stops attempting any self-describing agent installation for this product code until some action is taken by the administrator to correct the error. This correction might involve IBM Software Support. In this scenario, you must use the `tacmd deleteappinstallrec` command to clear the self-describing agent error record after you have addressed the problem. For more information about the `tacmd deleteappinstallrec` command, see the *IBM Tivoli Monitoring Command Reference*.

To determine if a self-describing agent product installation has failed with a terminal error condition, run the `tacmd listappinstallrecs` command by using the `-e` option to show error records only. For any error records with a STATE value of ME, the installation is not retried.

Retrying the self-describing agent installation

Complete the steps in this section to retry the self-describing agent installation.

Procedure

1. To avoid the same failure occurring again, you must first correct the condition that caused the installation to fail. The failure code is reported in the STATUS column of the `listappinstallrecs` output. In addition, the Tivoli Enterprise Monitoring Server message facilities (Audit, MSG2, and RAS1 messages) provide more information about the cause of the failure. Take corrective action to fix the condition or contact IBM Software support for assistance.
2. For each monitoring server, delete the failed installation records in the application properties table by running the `tacmd deleteappinstallrecs` command. Use this command to remove the blocking self-describing agent product installation record. See the *IBM Tivoli Monitoring Command Reference* for more information about using the `tacmd deleteappinstallrecs` command.
3. When each Tivoli Enterprise Monitoring Server failed product install record has been cleared, the monitoring server self-describing agent facility immediately notifies any running self-describing agent that can provide this level of product support, to retry the product installation. For example, if the previous install attempt for product *pc* and version 06230000 failed with a STATE of ME, and you run the `deleteappinstallrecs` command, any running *pc* agent for version 06230000 immediately retries the installation.
4. Run the `tacmd listappinstallrecs -t <pc>` again for product *pc* to determine the current installation state.

Enabling or disabling self-describing agent capability

The self-describing agent environment variable is set to Y to enable the self-describing agent capability at all components by default, with the exception of the hub monitoring server. The environment variable `KMS_SDA` is disabled (`KMS_SDA=N`) by default only on the hub monitoring server. All components that

connect to the hub monitoring server adjust their self-describing agent function to disabled if the hub monitoring server has the self-describing agent function disabled, or if the hub monitoring server is older and does not support self-describing agents.

This allows control of all self-describing agent enablement from a single point - the hub monitoring server. For more information, see “Enabling self-describing agent capability at the hub monitoring server” on page 204. A best practice is to only use the KMS_SDA variable at the hub monitoring server to control self-describing agent enablement, but it can be controlled at each component:

- Tivoli Enterprise Monitoring Server uses the KMS_SDA environment variable.
- Tivoli Enterprise Portal Server uses the TEPS_SDA environment variable.
- Self-describing agents use the TEMA_SDA environment variable.

You can also use `tacmd` commands to enable or disable the state of the self-describing agent function at the hub monitoring server, without having to recycle any hub monitoring servers. This control is only applied when the hub Tivoli Enterprise Monitoring Server is configured for self-describing agent capability using the environment variable `KMS_SDA=Y`. For more information, see “Dynamically controlling the hub monitoring server self-describing agent capability” on page 205.

For more information on self-describing agent variables, see “Environment variables” on page 893.

Chapter 12. Monitoring your operating system via a System Monitor Agent

A special type of Tivoli Management Services agent, the System Monitor Agent, allows you to send OS monitoring data directly to any SNMP or EIF event receiver such as Tivoli Netcool/OMNIBus, Tivoli Enterprise Console, or Tivoli NetView® without first passing the data to a Tivoli Enterprise Monitoring Server.

These lighter-weight agents (they require a much smaller footprint than full-function Tivoli Monitoring OS agents) are configured locally on the agent node, which makes them ideal for customers who want fully autonomous monitoring of their desktop operating systems in environments where disk space or image transmission bandwidth is in short supply. This configuration enables them to be deployed autonomously (that is, they do not require the support of a Tivoli Enterprise Monitoring Server): they send SNMP or EIF event information directly to an Event Collector such as IBM Tivoli Netcool/OMNIBus.

These agents can run in agent-only environments that lack the standard Tivoli Monitoring servers (the Tivoli Enterprise Monitoring Server and the Tivoli Enterprise Portal Server). Local configuration files can be placed on the system directly, or the System Monitoring Agent can be configured to check in on a schedule with a Central Configuration Server to collect updated configuration files. (See *Central Configuration Servers* in the *IBM Tivoli Monitoring Administrator's Guide*.)

You can also install or deploy System Monitor Agents using a Golden Master Image. To install System Monitor Agents using a Golden Master Image, use the following procedures.



To create the Golden Master Image:

1. Install a System Monitor Agent on a sample system.
2. Configure the agent as required:
 - You can provide the private configuration files, or provide the Central Configuration Connection information and test the image.
 - Once it is performing as desired, stop the agent.
3. Compress the files and directories within the IBM Tivoli Monitoring installation directory to create a Golden Master Image.

To install the Golden Master Image:

1. Transfer the compressed file to the target system.
2. Extract the image into the same installation directory.

Note: The agent installation directory must be the same on the original image and the target system.

3. Once the files are in place run the initialization script and start the agent:
 -  `<instdir>\InstallITM\UpdateAutoRun.cmd`
 -  `<instdir>/bin/UpdateAutoRun.sh`

These agents, which run on Windows and on Linux/UNIX, effectively replace the OMNIBus System Service Monitors for monitoring of desktop operating systems. IBM Netcool System Service Monitor users may find these agents appropriate as an upgrade path.

Note: Only agents created with version 6.2.2 (or subsequent) of the Agent Builder tool that do not include a data provider that requires Java can be installed with a System Monitor Agent. No other agents should be installed with a System Monitor Agent, nor should they be installed on nodes where there is already a Tivoli Monitoring component in operation.

For System Monitor Agents, you can use a silent installation process for both Windows and Linux/UNIX sites, along with operating system-specific configuration steps. There are also uninstallation processes for these Windows and Linux/UNIX OS agents. All these processes are detailed in this chapter.

“Background information about agent autonomy” on page 55 introduces the concept of agents that operate autonomously (that is, without a connection to a Tivoli Enterprise Monitoring Server). For additional information, see the *IBM Tivoli Monitoring Administrator's Guide*.

Installing the System Monitor Agent on Windows systems

The System Monitor Agent packages provide IBM Tivoli Monitoring OS agents whose reduced footprint makes them ideal for customers who need fully autonomous OS agent monitoring in environments where disk space or image transmission bandwidth is in short supply. These OS agent packages offer a significantly reduced bundle size and installed footprint.

The System Monitor Agents support direct Netcool/OMNIBus integration without a connection to a Tivoli Enterprise Monitoring Server. The System Monitor Agent instead sends SNMP alerts directly to an SNMP Event Collector such as Netcool/OMNIBus.

Installation times are significantly reduced due to the small size of the System Monitor Agents. These agents use simplified flat-file, machine-independent configuration. Configuration files no longer contain references to the host name or system where the file is located; with this change, you can share and transport common configuration files across your network.

Note:

1. Beginning with V6.3, Microsoft Visual C++ 2008 Redistributable Package (32 bit and 64 bit) is a prerequisite for the System Monitor Agent on Windows systems:

- Microsoft Visual C++ 2008 SP1 Redistributable Package (x86)
- Microsoft Visual C++ 2008 SP1 Redistributable Package (x64)

As both 32 bit and 64 bit agents can be installed on a 64 bit system, you should install both packages on a 64 bit Windows system. For more information go to the Microsoft Download Center at <http://www.microsoft.com/en-us/download/default.aspx>.

2. A System Monitor Agent must not be installed into a system where existing IBM Tivoli Monitoring components (including other monitoring agents) are already installed, with this exception: Agents built with Agent Builder V6.2.2 or subsequent may be installed alongside a System Monitor Agent, provided they run in the same mode as the Windows System Monitor Agent itself: if the Windows agent runs in 32-bit mode, only 32-bit Agent Builder agents are supported; if the Windows agent runs in 64-bit mode, only 64-bit Agent Builder agents are supported. 32-bit Agent Builder agents can be regenerated using Agent Builder to create 64-bit Windows binaries that you can install with the 64-bit Windows System Monitor Agent.
3. The NT Micro Agent in the `system_monitor_agent` directory of the IBM Tivoli Monitoring installation media only supports 8.3 format path values, with no special characters and no spaces, as an installation path. The bat files that install the NT Micro Agent do not support long paths or paths that are enclosed in double quotes. Use a short path without special characters or spaces to install the NT Micro Agent.

You must invoke a silent installation procedure to install the System Monitor Agent that monitors the Windows operating system:

```
silentInstall.cmd -p response_file
```

where:

response_file

is the name of the updated version of the nt_silent_install.txt response file you created that names the components to be installed and where.

On completion of the install script, the OS agent for Windows is installed, configured with a default configuration, and started.

Contents of the silent response file

The silent response file supports the following records:

comments

Any line beginning with a semicolon (;) is treated as a comment.

Keyword=value

These lines define variables in the silent response file. Do not include spaces before the keyword or immediately before or after the equal sign (=). Values must not contain the following characters:

\$
=
|

Note: By default the silentInstall.cmd installation script checks the processor architecture on which it is running; if it discovers it is running on a 64-bit x86_64 CPU, silentInstall.cmd installs the System Monitor Agent in 64-bit mode. Otherwise the agent is installed in 32-bit mode.

To force installation of the 32-bit System Monitor Agent on an x86_64 CPU, set variable **OS_ARCH** to 32 in the nt_silent_response.txt file. (**OS_ARCH** can also be set to 64, to force installation of a 64-bit agent. If **OS_ARCH** is blank or unspecified, the installer determines which agent to install based on the host processor.)

The silent response file supports the following variables:

License Agreement

(Required) This variable accepts the license agreement for the System Monitor Agent. In the sample response file, simply uncomment the line containing the **License Agreement** variable, or set the variable to **I agree to use the software only in accordance with the installed license**.

Install Directory

(Required) This variable specifies the directory where the System Monitor Agent will be installed. Characters supported for the **Install Directory** are alphanumeric characters, underscore (_), and hyphen (-).

KDC_FAMILIES_OVERRIDE

(Optional) This variable overrides the value assigned to the **KDC_FAMILIES** variable the agent uses at run time. In the sample response file, uncomment the line containing the **KDC_FAMILIES_OVERRIDE** variable and specify the wanted protocols to be enabled. If left commented out, the installer uses a value of:

IP.PIPE PORT:1918 IP use:n IP.SPIPE use:n SNA use:n

This example sets the default protocol to IP.SPIPE, using listening port 3660, and with all other protocols disabled:

KDC_FAMILIES_OVERRIDE=IP.SPIPE PORT:3660 IP.PIPE use:n IP use:n SNA use:n

For more information about the supported protocols and modifiers, see “Tivoli Monitoring protocol usage and protocol modifiers” on page 435.

Note: Failure to specify at least one active protocol results in the System Monitor Agent not starting. To rectify this you can manually edit the KDC_FAMILIES registry entry value.

EncryptionKey

(Optional) This variable sets the encryption key used for storing passwords and other protected content. The default value if not specified is **IBMTivoliMonitoringEncryptionKey**.

Note:

1. Do not use any of the following characters in your key:

&	ampersand
	pipe
'	single quote
=	equal sign
\$	dollar sign

In addition, do not specify double-byte (DBCS) characters.

2. Ensure that you document the value you use for the key. Use this key during the installation of any components that communicate with this monitoring server.

InstallLog

(Optional) This variable overrides the default location of the installation log created by the `silentInstall.cmd` command. The default value if not specified is `%CANDLE_HOME%\InstallITM\ibm_tivoli_micro_install.log`. In case of an installation error where the installation directory was not created, the installation log will be created in the `%TEMP%` directory as `ibm_tivoli_micro_install.log`. Characters supported for the **InstallLog** are alphanumeric characters, underscore (`_`), and hyphen (`-`).

A sample `nt_silent_install.txt` response file can be found in directory `\agents\system_monitor_agent\WINDOWS\` on the Agents DVD.

Configuring the System Monitor Agents on Windows

Environmental configuration for the autonomous OS agents is unchanged on Windows: Configuration parameters are contained in the agent's ENV file in the `InstDir\tmaitm6` directory, where *InstDir* is your installation directory. By directly editing this file, you can modify existing variables within that file or append additional environment variables.

User-defined local configuration files

These configuration files are stored in a separate local configuration directory, `InstDir\localconfig\pc`, where *pc* is the agent's two-character product code. This segregation simplifies agent configuration through the use of flat files. It also segregates your user configurations, which preserves them during agent upgrade. They also allow redistribution of a configured agent to other systems while ensure their configurations remain identical.

To modify the default location of the operational configuration files, perform the following steps:

1. Edit the agent's ENV file (for example, `KNTENV` for the Windows OS agent).
2. Add or modify the `IRA_LOCALCONFIG_DIR` parameter to specify a different local configuration directory location.

After you have modified the ENV file and saved it, recycle the agent to implement your updated configuration.

Note: On Windows, there are new CLI commands to start and stop an agent:

```
%CANDLE_HOME%\InstallITM\itmcmd.cmd agent start nt
```

```
%CANDLE_HOME%\InstallITM\itmcmd.cmd agent stop nt
```


Private situation configuration files

If a correctly named private situation configuration file is present in the System Monitor Agent's local configuration directory when the agent is started, the agent will run the private situations that the file defines. The file must be named `pc_situations.xml`, where `pc` is the agent's two-character product code.

You can use the `IRA_PRIVATE_SITUATION_CONFIG` parameter to specify a different file name or location. This variable is resolved relative to the `IRA_LOCALCONFIG_DIR`, or you can specify a fully qualified file name.

For more information about the private situation configuration files, see the *IBM Tivoli Monitoring Administrator's Guide*.

SNMP trap configuration files

If a correctly named trap configuration file is present in the System Monitor Agent's local configuration directory when the agent is started, the agent will emit the SNMP alerts that the file defines. The file must be named `pc_trapcnfg.xml`, where `pc` is the agent's two-character product code.

You can use the `IRA_EVENT_EXPORT_SNMP_TRAP_CONFIG` parameter to specify a different file name or location. This variable is resolved relative to the `IRA_LOCALCONFIG_DIR`, or you can specify a fully qualified file name.

For more information about SNMP trap configuration files, see the *IBM Tivoli Monitoring Administrator's Guide*.

Uninstalling the Windows System Monitor Agent

There is a separate procedure for uninstalling the System Monitor Agent that monitors the Windows operating system and deleting the directory where it is stored.

Issue the following command:

```
%CANDLE_HOME%\Install\ITM\uninstall.cmd [-f]
```

You are prompted to confirm your request to uninstall the agent. Respond **Y** to continue the uninstallation.

The `-f` option forces the uninstallation command to bypass the confirmation prompt.

If the command is invoked while the current working directory is `%CANDLE_HOME%\Install\ITM`, the directory is not deleted; you must delete it manually.

Note:

1. Uninstalling the Windows operating system agent also uninstalls all Agent Builder agents installed in the same environment.
2. Directories are not removed on Microsoft Windows systems if the command that attempts to remove it is running from the directory being removed or if a file is locked within that directory. Thus, if you run `uninstall.cmd` from `%CANDLE_HOME%\BIN`, the directory is not removed; you must remove it yourself.

Before running `uninstall.cmd`, it is recommended that you first close all processes (such as command prompts and text editors) currently accessing subdirectories of `%CANDLE_HOME%`. Then run the `uninstall.cmd` command from outside of `%CANDLE_HOME%`. Specify a fully qualified path.

If `uninstall.cmd` cannot remove a subdirectory, it displays the following message:

directory may have to be removed manually after this script completes.

This same message also is written to the uninstallation log file:
%TEMP%\IBM_Tivoli_Monitoring_System_Monitor_Agent_uninstall.log

Installing the System Monitor Agent on Linux or UNIX systems

The System Monitor Agent packages provide IBM Tivoli Monitoring operating-system agents whose reduced footprint makes them ideal for customers who need fully autonomous OS agent monitoring in environments where disk space or image transmission bandwidth is in short supply. These System Monitor Agent packages offer a significantly reduced bundle size and installed footprint.

The System Monitor Agents support direct Netcool/OMNIBus integration without a connection to a Tivoli Enterprise Monitoring Server. The System Monitor Agent instead sends SNMP alerts directly to an SNMP Event Collector such as Netcool/OMNIBus.

Installation times are significantly reduced due to the small size of the agents. System monitor agents use simplified flat-file, machine-independent configuration. Configuration files no longer contain references to the host name or system where the file is located; with this change, you can share and transport common configuration files across your network.

Note: A System Monitor Agent must not be installed into a system where existing IBM Tivoli Monitoring components (including other monitoring agents) are already installed, with this exception: Agent Builder agents built with Agent Builder V6.2.2 or subsequent may be installed alongside a System Monitor Agent.

There is a silent installation procedure that you must invoke to install these autonomous agents that monitor the Linux and UNIX operating systems:

```
silentInstall.sh -h installation_dir -p response_file
```

where:

installation_dir

is the directory on the target machine where the System Monitor Agent is to be installed.

response_file

is the name of the response file you created that names the components to be installed and where.

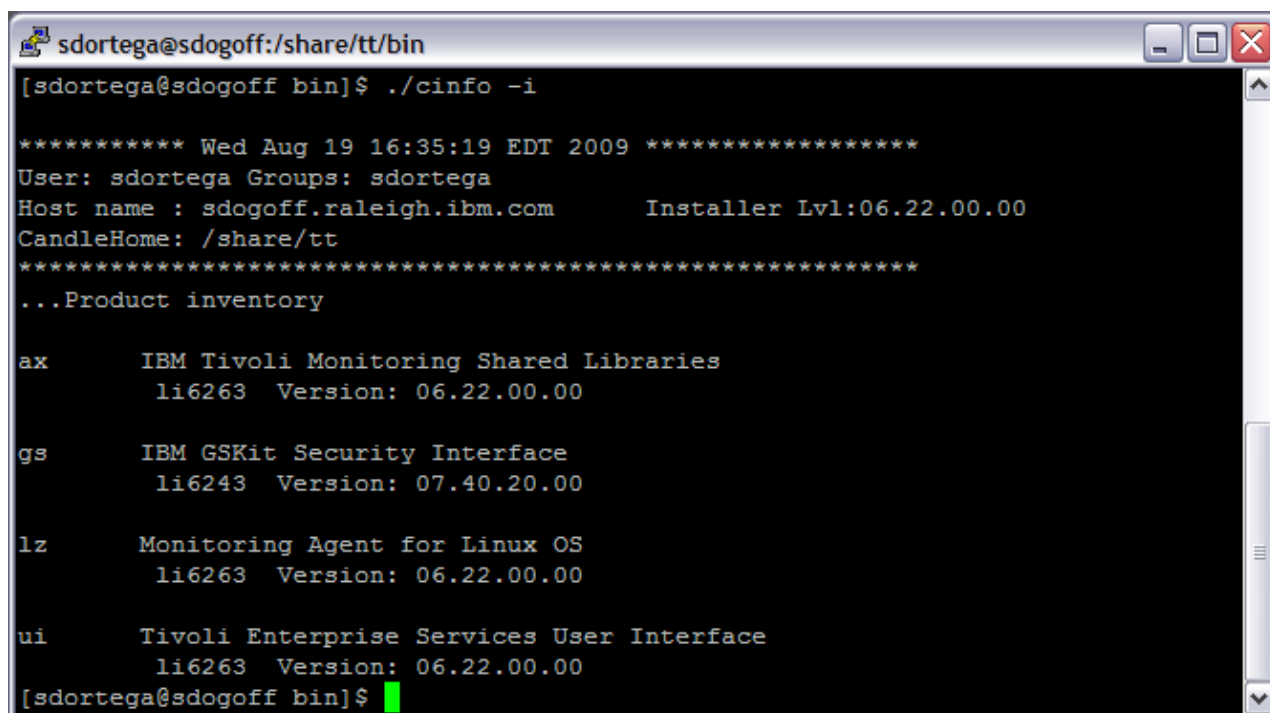
Note: You must not specify the path of the directory containing `./install.sh` as your IBM Tivoli Monitoring home directory. On certain platforms, this can cause the plug-in JAR files to overwrite themselves and become zero length files. The installation will fail as a result.

In addition, do not use the number sign (#) or exclamation mark (!) characters anywhere in your installation directory path.

To verify the installation, enter this command:

```
InstDir/bin/cinfo -i
```

where *InstDir* is the directory where you installed the System Monitor Agent. The list of installed agent components is displayed, as shown in Figure 84 on page 405.



```
sdortega@sdogoff:/share/tt/bin
[sdortega@sdogoff bin]$ ./cinfo -i

***** Wed Aug 19 16:35:19 EDT 2009 *****
User: sdortega Groups: sdortega
Host name : sdogoff.raleigh.ibm.com      Installer Lvl:06.22.00.00
CandleHome: /share/tt
*****
...Product inventory

ax      IBM Tivoli Monitoring Shared Libraries
        li6263  Version: 06.22.00.00

gs      IBM GSKit Security Interface
        li6243  Version: 07.40.20.00

lz      Monitoring Agent for Linux OS
        li6263  Version: 06.22.00.00

ui      Tivoli Enterprise Services User Interface
        li6263  Version: 06.22.00.00
[sdortega@sdogoff bin]$
```

Figure 84. Output of the `cinfo` command

If no agents are listed, check the installation logs for more information.

Contents of the silent response file

This section explains the contents of the silent response file.

The silent response file supports the following records:

comments

Any line beginning with a number sign (#) is treated as a comment.

Keyword=value

These lines define variables in the silent response file. Do not include spaces before the keyword or immediately before or after the equal sign (=). Values must not contain the following characters:

\$
=
|

The silent response file supports the following variables:

License_Agreement

(Required) This variable accepts the license agreement for the System Monitor Agent. In the sample response file, simply uncomment the line containing the **License Agreement** variable, or set the variable to **I agree to use the software only in accordance with the installed license**.

INSTALL_PRODUCT

(Required) Identifies the product you want installed; specify **lz** for Linux for **ux** for UNIX.

INSTALL_FOR_PLATFORM

(Optional) Identifies the platform for which the product should be installed; see Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883. If left commented out, the installer will use the platform of the machine on which the installation is performed.

INSTALL_ENCRYPTION_KEY

(Optional) This variable sets the encryption key used for storing passwords and other protected content. The default value if not specified is **IBMTivoliMonitoringEncryptionKey**.

Note:

1. Do not use any of the following characters in your key:

&	ampersand
	pipe
'	single quote
=	equal sign
\$	dollar sign

In addition, do not specify double-byte (DBCS) characters.

2. Ensure that you document the value you use for the key. Use this key during the installation of any components that communicate with this monitoring server.

INSTALL_IGNORE_RUNNING_PROCESSES

If set to **n**, this keyword aborts the installation of running IBM Tivoli Monitoring processes. If this parameter is set to **y**, the installation stops the running processes, installs the requested product, and restarts the processes that were running.

KDC_FAMILIES_OVERRIDE

(Optional) This variable overrides the value assigned to the **KDC_FAMILIES** variable that the agent uses at run time. In the sample response file, uncomment the line containing the **KDC_FAMILIES_OVERRIDE** variable and specify the wanted protocols to be enabled. If left commented out, the installer uses a value of:

```
IP.PIPE PORT:1918 IP use:n IP.SPIPE use:n SNA use:n
```

This example sets the default protocol to IP.SPIPE, using listening port 3660, and with all other protocols disabled:

```
KDC_FAMILIES_OVERRIDE='IP.SPIPE PORT:3660 IP.PIPE use:n IP use:n SNA use:n'
```

For more information about the supported protocols and modifiers, see “Tivoli Monitoring protocol usage and protocol modifiers” on page 435.

Note: Failure to specify at least one active protocol results in the System Monitor Agent not starting. To rectify this you can manually edit the KDC_FAMILIES environment parameter in the configuration file of the agent.

Sample `lz_silent_install.txt` and `ux_silent_install.txt` response files can be found in directory `/agents/system_monitor_agent/unix/` on the Agents DVD.

Configuring the System Monitor Agents on Linux or UNIX

The autonomous OS agents do not use the `pc.ini` files (where `pc` is the product code) for configuration. Instead these System Monitor Agents follow a configuration model based on flat files. These files are designed to segregate and preserve your user configuration, which simplifies agent upgrade. They also allow redistribution of a configured agent to other systems while ensuring the original configuration is exactly duplicated.

The System Monitor Agents' environment files are stored in the `InstDir/config` directory, where `InstDir` is your installation directory. When the agent is installed, two files are added to this directory:

.global.environment

This file contains IBM Tivoli Monitoring-defined global environment settings that are available to all monitoring agents that use the same installation directory. *User modifications should not be made to this file*, as the changes may not be preserved in future upgrades. Instead, make your global user configuration changes in the `global.environment` file described below.

.pc.environment

Here *pc* is the two-character IBM Tivoli Monitoring product code (for example, *.ux.environment* is the product-provided UNIX System Monitor Agent environment configuration file, and *.lz.environment* is the Linux System Monitor Agent environment configuration file). These files contain product-defined, agent-specific environment settings that are available only to the *pc* agents running in the same installation directory. Again, *user modifications should not be made to this file*, as the changes may not be preserved in future upgrades. Instead, make your user agent configuration changes in the *pc.environment* file described below.

Values defined in the *pc.environment* files override those defined in the global environment files.

In the *InstDir/config* directory (where *InstDir* is your installation directory), you may create custom environment files that override the environment settings defined in the *.global.environment* and *.pc.environment* files. These user customizations are preserved during future product upgrades. These files are:

global.environment

This file is not created when the System Monitor Agent is installed but must be created by you:

1. Copy the IBM Tivoli Monitoring-supplied *.global.environment* file to *global.environment*.
2. Delete all entries that do not need to be customized.
3. Add any custom global configuration variables that should be available to all agents running in the same installation directory.

User modifications should be made to this file as these changes will never be overwritten by future upgrades. Values specified in the *global.environment* configuration file override those in both the *.global.environment* file and the *.pc.environment*.

pc.environment

This file is not created when the System Monitor Agent is installed but must be created by you:

1. Copy the IBM Tivoli Monitoring-supplied *.pc.environment* file to *pc.environment*.
2. Delete all entries that do not need to be customized.
3. Add any agent-specific configuration variables that should be available to all *pc* agents running in the same installation directory.

User modifications should be made to this file as these changes will never be overwritten by future upgrades. Values specified in the user *pc.environment* file override all others.

User-defined local configuration files

These configuration files are stored in a separate local configuration directory, *InstDir/localconfig/pc*, and include this parameter:

IRA_LOCALCONFIG_DIR

Specifies a different local configuration directory. You may want to set the *IRA_LOCALCONFIG_DIR* parameter in the *pc.environment* file so each agent will have its own unique *IRA_LOCALCONFIG_DIR* setting. Alternatively, setting the *IRA_LOCALCONFIG_DIR* in the *global.environment* file allows all agents to share the same *IRA_LOCALCONFIG_DIR* setting.

Private situation configuration files

If a correctly named private situation configuration file is present in the System Monitor Agent's local configuration directory when the agent is started, the agent will run the private situations that the file defines. The file must be named *pc_situations.xml*, where *pc* is the agent's two-character product code.

The *IRA_PRIVATE_SITUATION_CONFIG* environment variable may be used to specify a different file name or location. This variable is resolved relative to the *IRA_LOCALCONFIG_DIR* directory, or you may specify a fully qualified filename.

For more information about the private situation configuration files, see the *IBM Tivoli Monitoring Administrator's Guide*.

SNMP trap configuration files

If a correctly named trap configuration file is present in the System Monitor Agent's local configuration directory when the agent is started, the agent will emit the SNMP alerts that the file defines. The file must be named *pc_trapcnfg.xml*, where *pc* is the agent's two-character product code.

The **IRA_EVENT_EXPORT_SNMP_TRAP_CONFIG** environment variable may be used to specify a different file name or location. This variable is resolved relative to the **IRA_LOCALCONFIG_DIR** directory, or you may specify a fully qualified filename.

For more information about SNMP trap configuration files, see the *IBM Tivoli Monitoring Administrator's Guide*.

Uninstalling the Linux or UNIX System Monitor Agent

This section outlines the procedure for uninstalling the System Monitor Agent on Linux or UNIX.

About this task

There is a separate, manual procedure for uninstalling the System Monitor Agent that monitors the Linux or UNIX operating system:

Procedure

1. Stop the agent by invoking either of the following commands:

- **Linux**

```
InstDir/bin/itmcmd agent stop lz
```

- **UNIX**

```
InstDir/bin/itmcmd agent stop ux
```

where *InstDir* is the directory where you installed the System Monitor Agent.

2. Stop any other agents running from the same *InstDir*.
3. Issue one of the following commands :

```
InstDir/bin/uninstall.sh
```

or:

```
InstDir/bin/uninstall.sh REMOVE EVERYTHING
```

Invoking the `uninstall.sh` script with the `REMOVE EVERYTHING` parameter removes all agent files and deletes the installation subdirectory tree.

Results

On UNIX and Linux, you can uninstall multiple, individual agents by listing each one in the **uninstall.sh** command.

Note: Uninstalling the UNIX or Linux OS System Monitor Agent also removes common files that are necessary for Agent Builder agents to function. If you are uninstalling the UNIX or Linux OS system monitor agent, you must also uninstall any Agent Builder agents in the same environment.

Defining common configuration parameters: accessing centralized configuration information

System Monitor Agents are started at the end of their silent installation. Unless private configuration files exist for them, these agents run but do not analyze situations or send events. Centralized configuration allows an agent to retrieve these files from the centralized configuration server and to begin using them immediately.

You define these parameters in the `pc_silent_install.txt` silent response file that is used when installing a System Monitor Agent.

The System Monitor Agent installation process uses entries in the silent response file to create entries in the agent's environment file, based on statements of this form: **SETENV***_parameter=value* statements create *parameter=value* statements in the agent's environment file, whereas **SETENCR***_parameter=value* create *parameter={AES256:keyfile:a}encryptedvalue* statements.

Example: These silent response file entries connect to a centralized configuration server running on the node named linuxhost:

```
SETENV_IRA_CONFIG_SERVER_URL= http://linuxhost:1920///linuxhost_lz/linuxhost_lz/  
SETENV_IRA_CONFIG_SERVER_USERID=itmuser  
SETENCR_IRA_CONFIG_SERVER_PASSWORD=password  
SETENV_IRA_CONFIG_SERVER_FILE_PATH=initloadlist/@PRODUCT@  
SETENV_IRA_CONFIG_SERVER_FILE_NAME=cnfglist.xml
```

Agents recognize the following keywords and substitute for them runtime values retrieved from the client:

@PRODUCT@

Agent's lowercase, two-letter product code. Example: For a Windows OS agent, **@PRODUCT@_trapcnfg.xml** resolves to **nt_trapcnfg.xml**.

@ITMHOME@

IBM Tivoli Monitoring installation path. Example: If this is a Linux system and the default installation path is used, **@ITMHOME@** resolves to **/opt/IBM/ITM/**.

@MSN@

Agent's managed system name (not the subnode name). Example: If the agent's managed system name is primary:icvw3d62:nt, **@MSN@** resolves to **primary-icvw3d62-nt**.

@TASKNAME@

Agent's process name. Examples: **klzagent**; **kntcma**.

@VERSION@

Agent's product version string. Example: If the agent's version is Tivoli Monitoring 6.2.2 fix pack 2, **@VERSION@** resolves to **06-22-02**.

@HOSTNAME@

Computer host name. Example: **myhost**.

@IPADDRESS@

Computer network interface IP address. Example: If the agent's IP address is 9.42.38.333, **@IPADDRESS@** resolves to **9-42-38-333**.

@OSTYPE@

Operating system type. Examples: **linux**; **win2003**.

@OSVERSION@

Operating system version. Examples: Red Hat Enterprise Linux Version 5 (64 bit) resolves to **2-6-18-128-el5**; Windows 2003 (32 bit) with Service Pack 2 resolves to **5-2-sp2**.

@SYSTEMID@

Computer system identifier. Example: System ID icvr4a04.mylab.mycity.ibm.com resolves to **icvr4a04-mylab-mycity-ibm-com**.

For detailed information on using the centralized configuration facility, see the *IBM Tivoli Monitoring Administrator's Guide*.

Note:

1. All special characters in the parameter values for all keywords other than @ITMHOME@ are converted to dashes (-). For example, if the IP address is 9.42.38.233, keyword @IPADDRESS@ resolves to **9-42-38-233**.

The value for @ITMHOME@, however, remains unchanged.

2. The value of SETENCR_IRA_CONFIG_SERVER_PASSWORD may be either plain text or encrypted when saved in the *pc_silent_install.txt* silent response file. Plain-text values are encrypted when created in the agent environment file. Encrypted values are created as specified. The *itmpwdsnmp* utility is used interactively on another system to encrypt the password string if desired; see the *IBM Tivoli Monitoring Administrator's Guide*.

Chapter 13. Configuring IBM Tivoli Monitoring components

Although the majority of configuration is done during the product installation, you can use the Manage Tivoli Enterprise Monitoring Services tool to configure components at any time.

You can also use the Manage Tivoli Enterprise Monitoring Services tool to start and stop components.

Note: You can also perform many of these configuration and start and stop procedures from the command-line. Where this is possible, the command is included. See the for a complete description, including parameters, of the commands that you can use in the installation and configuration of IBM Tivoli Monitoring.

You can perform the tasks in Table 67 with the Manage Tivoli Enterprise Monitoring Services tool:

Table 67. Configuration tasks available through Manage Tivoli Enterprise Monitoring Services

Task	Where to find information
Start Manage Tivoli Enterprise Monitoring Services	"Starting Manage Tivoli Enterprise Monitoring Services"
Change the configuration of the monitoring server	"Changing the configuration of the Tivoli Enterprise Monitoring Server" on page 412
Change the configuration of the Tivoli Enterprise Monitoring Automation Server	"Configuring or changing the Tivoli Enterprise Monitoring Automation Server connection" on page 414
Configure agents and other monitoring components	"Configuring or changing the monitoring server connection for agents" on page 417
Start and stop components	"Starting and stopping components" on page 420
Monitor the status of remote monitoring servers and agents by configuring heartbeat monitoring.	"Configuring the heartbeat interval" on page 426

Starting Manage Tivoli Enterprise Monitoring Services

Depending on the operating system you are using, the procedure for starting Manage Tivoli Enterprise Monitoring Services is different.

Follow the steps in this chapter to start Manage Tivoli Enterprise Monitoring Services on Windows or Linux/UNIX computers.

Starting Manage Tivoli Enterprise Monitoring Services on Windows computers

Use the following steps to start Manage Tivoli Enterprise Monitoring Services on a computer running Windows.

1. Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Monitoring Services**.

Starting Manage Tivoli Enterprise Monitoring Services on Linux or UNIX computers

Use the following steps to start Manage Tivoli Enterprise Monitoring Services on a computer running Linux or UNIX.

Procedure

1. Change to the bin directory:
`cd install_dir/bin`
2. Run the following command by using the parameters described in the next table:
`./itmcmd manage [-h install_dir] [-s]`

-h	(optional) An option used to specify the installation directory.
install_dir	The installation directory for IBM Tivoli Monitoring.
-s	<p>(optional) Option to specify safe mode operation.</p> <p>Safe mode invokes the JRE with the -nojit option (no just-in-time compiler). If you encounter a Java failure error, try running the command as before, but also specifying the -s option.</p> <p>Entering the above commands with -? displays the syntax for using the -s option.</p>

Results

The Manage Tivoli Enterprise Monitoring Services utility is displayed.

Note: Note that the **Platform** column for agents lists the platform that the binary code was built on, not the platform that you are running on.

Changing the configuration of the Tivoli Enterprise Monitoring Server

You can change the basic configuration of the monitoring server through Manage Tivoli Enterprise Monitoring Services.

Procedure

1. In the Manage Tivoli Enterprise Monitoring Services window, right-click the monitoring server.
2. Click **Reconfigure** (on Windows) or **Configure** (on UNIX).
3. Identify the communications protocol for the monitoring server. You have four choices: IP.UDP, IP.PIPE, IP.SPIPE, or SNA. You can specify three methods for communication - this enables you to set up backup communication methods. If the method you've identified as Protocol 1 fails, Protocol 2 is used.
4. Click **OK**.
5. Complete the fields settings listed in Table 68 for the communications protocol for the monitoring server and click **OK**.

Table 68. Communications protocol settings

Field	Description
IP.UDP Settings	
Hostname or IP Address	The host name or IP address for the hub monitoring server. Note that the Tivoli Enterprise Monitoring Server supports both IPV4 and IPV6 addressing formats.
Port # or Port Pools	The listening port for the hub monitoring server.
IP.PIPE Settings	

Table 68. Communications protocol settings (continued)

Field	Description
Hostname or IP Address	The host name or IP address for the hub monitoring server. Note that the Tivoli Enterprise Monitoring Server supports both IPV4 and IPV6 addressing formats.
Port Number	The listening port for the monitoring server. The default number is 1918.
IP.SPIPE Settings	
Hostname or IP Address	The host name or IP address for the hub monitoring server. Note that the Tivoli Enterprise Monitoring Server supports both IPV4 and IPV6 addressing formats.
Port number	The listening port for the hub monitoring server. The default value is 3660.
SNA Settings	
Network Name	The SNA network identifier for your location.
LU Name	The LU name for the monitoring server. This LU name corresponds to the Local LU Alias in your SNA communications software.
LU 6.2 LOGMODE	The name of the LU6.2 LOGMODE. The default value is "CANCTDCS."
TP Name	The transaction program name for the monitoring server.

- If the monitoring server was running when you began the configuration process, after the reconfiguration is complete, you are asked if you want it restarted.

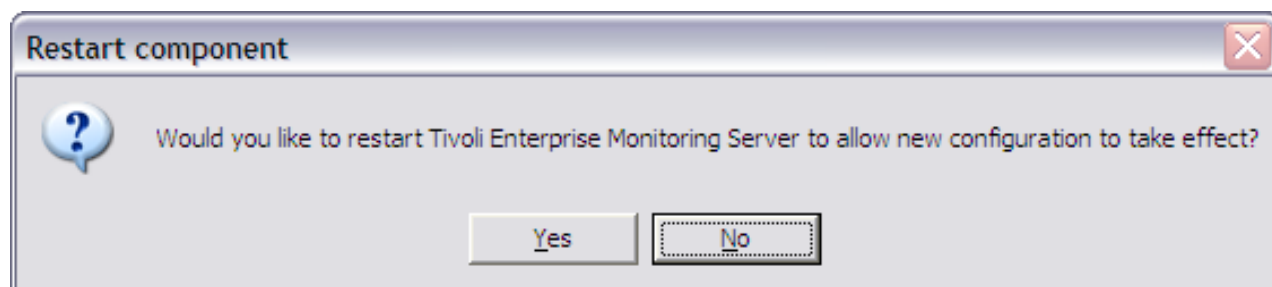


Figure 85. Restart Component window: Tivoli Enterprise Monitoring Server

Reply **Yes** or **No**.

What to do next

On Linux and UNIX, you can also use the **itmcmd config -S** command to change the configuration of a monitoring server. When the CLI completes the reconfiguration, if the monitoring server was running when you began this process, you are asked if you want it restarted:

Would you like to restart the component to allow new configuration to take effect [1=Yes, 2=No]
(Default is: 2):

Reply **1** or **2**, as appropriate.

If you choose the restart, the monitoring server is stopped and then started again. These actions are necessary to force the server to read your changed configuration (which is always read at server startup).

On UNIX platforms the component should be restarted with the same user that it previously ran on. If the monitoring server was not running when reconfigured, no action is performed, and the server remains stopped.

Note:

1. Use caution when starting, stopping, or restarting the Tivoli Enterprise Monitoring Server, as it is a key component.
2. You cannot configure a hub monitoring server and a remote monitoring server on the same system because you cannot have two processes listen to the same IP port on the same system.
 - You can configure any number of remote monitoring servers on the same system as long as each reports to a different hub and uses a different port number.
 - You can configure any number of hub monitoring servers on the same system as long as each uses a different port number.
 - If a hub monitoring server and a remote monitoring server are configured on the same system, the remote monitoring server must report to a hub on another system using a port other than the one used by the hub running on the same system.
 - You cannot have two monitoring servers talking to each other over IP from the same system unless one of them is a high-availability hub monitoring server, because a high-availability hub is isolated to a private IP address. See the *IBM Tivoli Monitoring High Availability Guide for Distributed Systems*.

Configuring or changing the Tivoli Enterprise Monitoring Automation Server connection

You must configure the Tivoli Enterprise Monitoring Automation Server if you are using the Registry Services component of IBM Jazz for Service Management and other products to integrate with IBM Tivoli Monitoring using OSLC linked data interfaces.

Before configuration, you must be aware of the existing Registry Services configuration details. See “Required software and memory requirements for using the OSLC Performance Monitoring service provider” on page 94 for more information on setting up Registry Services and optionally Security Services before installing the monitoring automation server.

The monitoring automation server cannot be used if the Hub monitoring server is configured for the Hot Standby high availability feature.

Windows: Configuring a Tivoli Enterprise Monitoring Automation Server

You can change the basic configuration of the automation server through Manage Tivoli Enterprise Monitoring Services.

Before you begin

A locally installed hub monitoring server is a prerequisite for the Tivoli Enterprise Monitoring Automation Server. An error message reports a discrepancy if you attempt to configure the automation server with a remote monitoring server. You must configure the automation server connection information before you can start it.

Procedure

1. In the Manage Tivoli Enterprise Monitoring Services window, right-click the Tivoli Enterprise Monitoring Automation Server.
2. Click **Reconfigure**.

3. Complete your Registry Services configuration details and Tivoli Enterprise Monitoring Automation Server connection details. When you select **Configure Registry Services connection**, you must provide the HTTP protocol, hostname, and port for clients to connect to the automation server. The **User ID** and **Password** fields must be completed with the User ID and Password values for your Registry Services provider registry.
In the **Automation Server connection information** section, provide the HTTP protocol and port for clients to connect to the automation server.

The screenshot shows a Windows-style dialog box titled "Configure Tivoli Enterprise Monitoring Automation Server". It contains two main sections. The first section, "Registry Services connection information", has a checked checkbox "Configure Registry Services connection". Below it, the "Protocol" is set to "HTTP" (selected with a radio button), and "HTTPS" is unselected. The "Hostname" field contains "server", "Port number" contains "16310", "User ID" contains "system", and both "Password" and "Confirm password" fields are filled with asterisks. The second section, "Automation Server connection information", has "Protocol" set to "HTTP" and "Port number" set to "10001". At the bottom of the dialog are three buttons: "OK", "Cancel", and "Help".

Figure 86. Configure Tivoli Enterprise Monitoring Automation Server

4. Click OK

Results

The Manage Tivoli Enterprise Monitoring Services utility is opened. (This might take a few minutes.)

What to do next

To configure the Performance Monitoring service provider to use the Security Services component of Jazz for Service Management to authenticate OSLC client applications, you must set the `KAS_SECURITY_SERVICES_ENABLED` environment variable to Yes in the Tivoli Enterprise Monitoring Automation Server's environment file and restart the monitoring automation server.

Linux/UNIX: Configuring a Tivoli Enterprise Monitoring Automation Server

Choose a procedure for configuring the automation server on Linux/UNIX.

Use one of the following procedures to configure the Tivoli Enterprise Monitoring Automation Server:

- "Command-line procedure"
- "GUI procedure" on page 416

Command-line procedure

Complete the steps in this section to configure the Tivoli Enterprise Monitoring Automation Server, by using the Linux or UNIX command-line.

Procedure

1. At the command-line, change to the /opt/IBM/ITM/bin directory (or the directory where you installed IBM Tivoli Monitoring).
2. Run the following command:
`./itmcmd config -A as`

where **as** is the product code of the Tivoli Enterprise Monitoring Automation Server.

3. Enter your Registry Services connection settings. When you set the Edit Registry Services connection settings parameter to Yes, you must specify the correct Protocol, Hostname, and Port number for the WebSphere server where Registry Services is installed. You must also enter a user ID and Password for the Registry Services.

This section contains an example of the command-line syntax. In this example, you must replace *server* with the host name for your Registry Services provider registry:

```
Edit Registry Services connection settings [1=Yes, 2=No] (Default is: 1):
Protocol [1=HTTP, 2=HTTPS] (Default is: 1): 1
Hostname(Default is: ): server
HTTP Port Number(Default is: 16310):
Set Registry Services user and password [1=Yes, 2=No] (Default is: 1):
User(Default is: ): user
Password:
Re-type: Password:
```

4. Provide the HTTP protocol and port for clients to connect to the automation server.
Automation Server connection information
Protocol [1=HTTP, 2=HTTPS] (Default is: 1):
HTTP Port Number(Default is: 10001):

What to do next

To configure the Performance Monitoring service provider to use the Security Services component of Jazz for Service Management to authenticate OSLC client applications, you must set the KAS_SECURITY_SERVICES_ENABLED environment variable to Yes in the Tivoli Enterprise Monitoring Automation Server's environment file and restart the monitoring automation server.

GUI procedure

Complete the steps in this section to use the Manage Tivoli Enterprise Monitoring Services window on a Linux Intel or UNIX computer to configure the Tivoli Enterprise Monitoring Automation Server.

Procedure

1. Log on to the computer where the automation server is installed.
2. Start the Manage Tivoli Enterprise Monitoring Services utility:
 - a. Change to the bin directory:
`cd install_dir/bin`
 - b. Run the following command by using the parameters described in the next table:
`./itmcmd manage [-h ITMinstall_dir]`

where:

-h	(optional) An option used to specify the installation directory.
ITMinstall_dir	The directory where the monitoring server is installed. The default installation directory is /opt/IBM/ITM.

The Manage Tivoli Enterprise Monitoring Services window is displayed.

3. Right-click **Tivoli Enterprise Monitoring Automation Server** and click **Configure**.

4. Complete your Registry Services connection details and automation server connection details. When you select **Configure Registry Services connection**, you must specify the correct Protocol, Hostname, and Port number for the WebSphere server where Registry Services is installed.

The **User ID** and **Password** fields must be completed with the User ID and Password values for your Registry Services provider registry. Click **OK** to continue.

In the **Automation Server connection information** section, provide the HTTP protocol and port for clients to connect to the automation server.

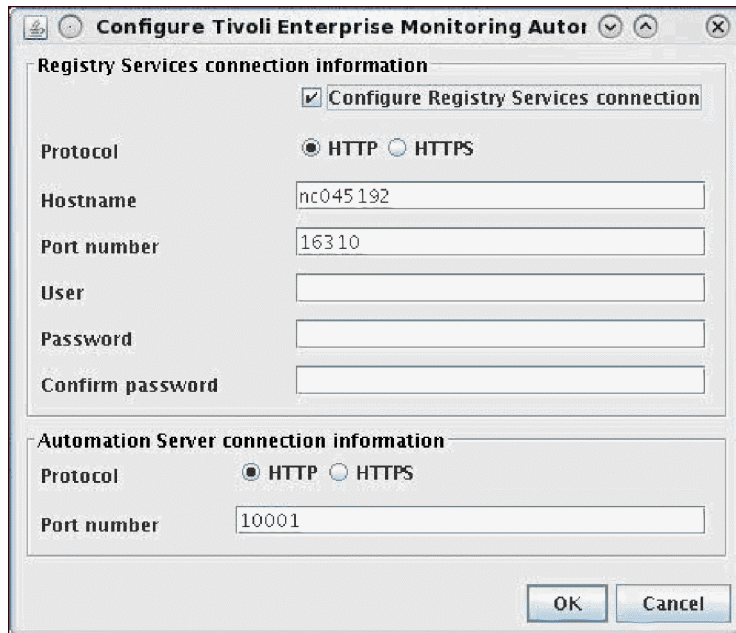


Figure 87. Configure Tivoli Enterprise Monitoring Automation Server

5. If the automation server is already running, after reconfiguring it you are prompted whether you want to restart the automation server or not. You can also stop and restart the automation server from the Manage Tivoli Enterprise Monitoring Services window:
 - a. Right-click **Tivoli Enterprise Monitoring Automation Server** and click **Stop**.
 - b. Right-click **Tivoli Enterprise Monitoring Automation Server** and click **Start**.

What to do next

To configure the Performance Monitoring service provider to use the Security Services component of Jazz for Service Management to authenticate OSLC client applications, you must set the `KAS_SECURITY_SERVICES_ENABLED` environment variable to Yes in the Tivoli Enterprise Monitoring Automation Server's environment file and restart the monitoring automation server.

Configuring or changing the monitoring server connection for agents

Use this procedure to configure or change the monitoring server connection to the agents.

Procedure

1. In the Manage Tivoli Enterprise Monitoring Services window, select the agent whose connection you want to configure. You can select multiple agents by holding down the Shift key or Control key and selecting agents.
2. Click **Actions** → **Reconfigure**.

3. Identify the communications protocol for communication with the monitoring server. You have four choices: IP.UDP, IP.PIPE, IP.SPIPE, or SNA. You can specify three methods for communication - this enables you to set up backup communication methods. If the method you've identified as Protocol 1 fails, Protocol 2 is used.
4. Click **OK**.
5. Complete the following fields and click **OK**:

Table 69. Communications protocol settings

Field	Description
IP.UDP Settings	
Hostname or IP Address	The host name or IP address for the hub monitoring server. Note that the Tivoli Enterprise Monitoring Server supports both IPV4 and IPV6 addressing formats.
Port # or Port Pools	The listening port for the hub monitoring server.
IP.PIPE Settings	
Hostname or IP Address	The host name or IP address for the hub monitoring server. Note that the Tivoli Enterprise Monitoring Server supports both IPV4 and IPV6 addressing formats.
Port Number	The listening port for the monitoring server. The default number is 1918.
IP.SPIPE Settings	
Hostname or IP Address	The host name or IP address for the hub monitoring server. Note that the Tivoli Enterprise Monitoring Server supports both IPV4 and IPV6 addressing formats.
Port number	The listening port for the hub monitoring server. The default value is 3660.
SNA Settings	
Network Name	The SNA network identifier for your location.
LU Name	The LU name for the monitoring server. This LU name corresponds to the Local LU Alias in your SNA communications software.
LU 6.2 LOGMODE	The name of the LU6.2 LOGMODE. The default value is "CANCTDCS."
TP Name	The transaction program name for the monitoring server.

6. If the agent was running when you began the configuration process, after the reconfiguration is complete, you are asked if you want the agent restarted.




Figure 88. Restart of Monitoring Agent window

Reply **Yes** or **No**.

What to do next

On Linux and UNIX, you can also use the **itmcmd config -A** command to change the configuration of a monitoring agent. When the CLI completes the reconfiguration, if the agent was running when you began this process, you are asked if you want the agent restarted:

Would you like to restart the component to allow new configuration to take effect [1=Yes, 2=No]
(Default is: 1):

Reply **1** or **2**, as appropriate.

If you choose the restart, the agent is stopped and then started again. These actions are necessary to force the agent to read your changed configuration (which is always read at agent startup). On UNIX platforms the component should be restarted with the same user that it previously ran on. If the agent was not running when reconfigured, no action is performed, and the agent remains stopped.

Note: If you upgrade an agent and the upgrade includes new columns for an existing attribute group, you must stop and restart history collection to get the new attributes to be picked up by the Tivoli Enterprise Monitoring Server when history is being collected at the Tivoli Enterprise Monitoring Server. This applies to self-describing agents as well.

Configuring the handling of pure events by the monitoring server

When the Tivoli Enterprise Monitoring Server has EIF Event Integration enabled, and data arrives from agents at a rapid pace, multiple rows of data either get put in a single event going to the event receiver or they get discarded. You can use the following procedure to configure the monitoring server to support one event per row of pure data. This ensures that when multiple rows of data arrive from the agent in quick succession, they are not concatenated into a single event; each row of data will generate a corresponding event.

Procedure

1. Create a file KPXATRGP on the Tivoli Enterprise Monitoring Server in the following directory:
 - \$CANDLEHOME\CMS (C:\IBM\ITM\CMS) for Windows systems.
 - \$CANDLEHOME/tables/ for UNIX systems.

2. Add the entries for pure event tables in the format <application-name>.<table-name> followed by a new line, for example:

```
KNT.NTEVTLOG  
TES00.TES123450
```

3. Always restart the Tivoli Enterprise Monitoring Server after adding a new entry to this file.

Results

This enables the feature for pure event tables only. Each row of data from the agent will result in an event. You should verify that all events from the agents show up in the EIF cache file or Event Server. The feature cannot be enabled for sample events. If there is a constant high volume of data from the agents, enabling this feature for all pure event tables will slow down the Tivoli Enterprise Monitoring Server performance. If you want to enable this feature you must enable it on all Tivoli Enterprise Monitoring Servers (hub and remote) in the environment.

Starting and stopping components

You can start and stop the IBM Tivoli Monitoring components from Manage Tivoli Enterprise Monitoring Services.

Procedure

1. Right-click the component (such as a specific agent or the Tivoli Enterprise Portal Server) that you want to start or stop.
2. Click **Start**, **Stop**, or **Recycle** (Windows only) from the menu.

What to do next

You can also use the following commands to start and stop components, including agents built by the Agent Builder and System Monitor Agents:

itmcmd server

Starts and stops a UNIX monitoring server.

itmcmd agent

Starts and stops a UNIX monitoring agent.

tacmd startAgent

Starts both Windows, Linux, and UNIX monitoring agents.

tacmd stopAgent

Stops both Windows, Linux, and UNIX monitoring agents.

See the *IBM Tivoli Monitoring Command Reference* for the syntax of these commands.

Note: When a hub Tivoli Enterprise Monitoring Server is recycled, all current events are recycled.

Specifying network interfaces

If there are multiple TCP/IP interfaces on the computer on which a monitoring server is running, you need to identify which interfaces monitoring agents or the Tivoli Enterprise Portal Server should use when connecting to the monitoring server. Setting network interfaces affects all the components installed on the local computer.

Before you begin

For instructions on configuring a network interface list on z/OS, see *Configuring the Tivoli Enterprise Monitoring Server on z/OS* (<http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/>)

About this task

To specify the network interfaces to be used by the portal to connect to a hub monitoring server, or by a monitoring agent to connect to a hub or remote, complete these steps:

Procedure

1. In the Manage Tivoli Enterprise Monitoring Services window, select **Actions** → **Advanced** → **Set Network Interface**.
2. On the Set Desired Network Interface window specify the network interface or interface you want to use
Specify each network adapter by the host name or IP address to be used for input and output. Use a blank space to separate the entries. If your site supports DNS, you can specify IP addresses or short host names. If your site does not support DNS, you must specify fully qualified host names.
3. Click **OK** to close save the settings and close the window.

Controlling port number assignments

IBM Tivoli Monitoring assigns port numbers to each component in an installation to be used for communication with other components. Default well-known ports are reserved for major components such as monitoring servers and the portal server. For all other components, an algorithm calculates the listening port to reserve.

You might want to change or add to the default assignments under some conditions, for example, when the default port assigned to a monitoring agent has already been reserved by another application, or the portal server requires a second port number for communication through a firewall with Network Address Translation (NAT).

Configuring port number assignments for the monitoring server

The default IP.UDP and IP.PIPE listening port setting for the Tivoli Enterprise Monitoring Server is 1918. For IP.SPIPE, it is 3660. For SNA, it is 135.

While you can specify a different port during or after installation, it is best to use the default setting. To reconfigure the port after installation, see “Configuring or changing the monitoring server connection for agents” on page 417.

Configuring port number assignments for the monitoring automation server

The default port number that the Tivoli Enterprise Monitoring Automation Server listens on for incoming HTTP requests is 10001. The default port for listening for HTTPS incoming requests is 10002.

You can specify a different port during or after installation. To reconfigure the port after installation, see “Configuring or changing the Tivoli Enterprise Monitoring Automation Server connection” on page 414.

Configuring port number assignments for the portal server

Communications between Tivoli Enterprise Portal clients and the Tivoli Enterprise Portal Server are controlled by a portal server interface definition.

The default interface definition assigns port 15001 to the Tivoli Enterprise Portal Server and ports 15200 (for HTTP requests) and 15201 (for HTTPS) to the integrated Web server that is installed with the portal

server. You can define additional interfaces to allow access through a firewall with NAT or through a secondary Network Interface Card (NIC). See “Firewall network address translation (NAT) or multiple network interface cards” on page 876.

The following sections describe how to change port number assignments on the portal server for connections to Tivoli Enterprise Portal clients (browser clients and desktop clients).

Changing the port number for the IBM Tivoli Monitoring Service Console

A portal server on Windows, Linux, or UNIX uses port 15200 for HTTP connections and 15201 for HTTPS connections.

Note: Ports 1920 and 3361 are legacy ports that are needed for connection to the IBM Tivoli Monitoring Service Console. However, these ports should no longer be used for clients connecting to the portal server. Ports 15200 and 15201 replace ports 1920 and 3361 for clients of the portal server.

Do not change the default port settings, especially on multifunction UNIX and Linux systems, since many components might be located on the same system and some of these components might depend on the default values being used for HTTP and HTTPS ports.

If you need to change the default settings, you can change them by using the KDE_TRANSPORT environment variable:

Windows

1. In the Manage Tivoli Enterprise Monitoring Services window, right-click **Tivoli Enterprise Portal Server**, point to **Advanced**, and select **Edit ENV File** to open the KFWENV file.
2. Add the following line to the file:
`KDE_TRANSPORT=HTTP:1920 HTTPS:3661`
Substitute the port numbers you want to use.
3. If a KDC_FAMILIES environment variable exists in the file, copy the settings from that variable to KDE_TRANSPORT (except the ones you want to override). KDE_TRANSPORT supersedes and overrides KDC_FAMILIES.
4. Save the file.
5. Recycle the portal server. (Right-click **Tivoli Enterprise Portal Server** and select **Recycle**.)

Linux

AIX

1. Change to the *install_dir*/config directory (where *install_dir* is the IBM Tivoli Monitoring installation directory).
2. Add the following line to the cq.ini file:
`KDE_TRANSPORT=HTTP:1920 HTTPS:3661`
Substitute the port numbers you want to use.
3. If a KDC_FAMILIES environment variable exists in the file, copy the settings from that variable to KDE_TRANSPORT (except the ones you want to override). KDE_TRANSPORT supersedes and overrides KDC_FAMILIES.
4. Recycle the portal server.

Special settings:

The KDE_TRANSPORT environment variable keyword redefines the ports to be used by the HTTP and HTTPS daemons. Certain special settings are also provided to meet special needs.

HTTP:0

Forces all access to the integrated web server to use the secure HTTPS daemon service.

HTTPS:0

If SSL is not available on the local system, this KDE_TRANSPORT environment variable eliminates startup error messages pertaining to a failure to bind an HTTPS (3661) listening port.

HTTP:0 HTTPS:0

These settings block access to both the HTTP (1920) and HTTPS (3661) listening ports, which has the effect of disabling all usage of the HTTP and HTTPS daemon services.

HTTP_SERVER:n

This KDE_TRANSPORT environment variable keyword disables HTTP and HTTPS daemon services. Do not specify this for a hub monitoring server or for the portal server.

HTTP_CONSOLE:n

This KDE_TRANSPORT environment variable keyword disables the CT/Service Console facility of the HTTP daemon service. HTTP_CONSOLE:N removes the process from the published Tivoli service index; this makes the process inaccessible from the CT/Service Console.

Changing the port number for desktop client connections to the portal server

Use the following procedure for each desktop client instance you want to change.

Procedure

1. In the Manage Tivoli Enterprise Monitoring Services window, right-click the Tivoli Enterprise Portal Desktop instance that you want to change, and select **Reconfigure**.

The Configure Application Instance window opens.

2. In the **Parms** list, scroll down to `cnp.http.url.port` and double-click.

3. On the Edit window, perform the following steps:

- a. Change the port number value to the port number you want.
- b. Select the **In Use** check box.

Note: If you fail to select this check box, the port number value that you entered will revert to the original value.

- c. Click **OK**.
4. On the Configure Application Instance window, verify that the port number value (in the **Value** column) for `cnp.http.url.port` has changed.
 5. Click **OK**.

Configuring port number assignments for monitoring agents

If multiple Tivoli Monitoring agents are installed on the same system, they cannot share the same port number for IP.PIPE or IP.SPIPE communications. An automatic mechanism has been developed to allocate additional ports for communications to the monitoring server. IBM Tivoli Monitoring uses the following algorithm to allocate port numbers for monitoring agents to reach the monitoring server: *reserved port = well-known port + (N*4096)*

where:

- *well-known port* is the port number assigned to the monitoring server, for example, 1918.
- *N* indicates the position of the monitoring agent in the startup sequence for agents.

For example, if there are two monitoring agents on a system, and the monitoring server uses port 1918, the first monitoring agent in the startup sequence is assigned port 6014 ($1918 + 1*4096$) and the second agent to start is assigned port 10110 ($1918 + 2*4096$). These port assignments occur automatically to avoid conflicts between multiple monitoring agents.

For “piped” protocols such as IP.PIPE and IP.SPIPE (but not IP or SNA), you also have the ability to control how port numbers are assigned to a monitoring agent by using the SKIP and COUNT parameters

on the KDE_TRANSPORT environment variable. These two parameters are described in more detail below, but here is an example of how you might use them when configuring a monitoring agent:

```
KDE_TRANSPORT=IP.PIPE PORT:1918 COUNT:1 SKIP:2 IP use:n SNA use:n IP.SPIPE use:n
```

See also the following information about using the IP.PIPE and IP.SPIPE protocols and parameters:

- The PORT parameter specifies the well-known port for the monitoring server.
- The piped protocols allocate one, and only one, physical port per process.
- There can be at most 16 IP.PIPE/IP.SPIPE processes per host.
- If you need to run more than 16 processes on a host, you can choose one of the following options:
 1. Use IP.UDP (User Datagram Protocol) for connections except to and from the Tivoli Enterprise Portal Server. The use of UDP is not recommended for the portal server.
 2. Add EPHEMERAL:Y as an optional KDE_TRANSPORT setting, for example:

```
KDE_TRANSPORT=IP.PIPE PORT:1918 EPHEMERAL:Y IP use:n SNA use:n
```

Marking an agent as ephemeral means that it will not try to open its own listening port. Instead, it will only initiate one-way, outbound communication to the monitoring server.

- The COUNT:N parameter is the mechanism for reserving IP.PIPE ports for components that connect to the monitoring server. N is the number of IP.PIPE ports to reserve on a host in addition to the well-known port for the monitoring server. Use the COUNT parameter to reserve ports for components that must be accessible from outside a firewall. Accessibility from outside the firewall requires IP.PIPE ports and because these ports must be permitted at the firewall, the ports must be predictable.

For example, if the well-known port is 1918, COUNT:3 starts the search at port 6014 ($1918 + 1 \times 4096$). If the agent process cannot bind to port 6014, the algorithm tries port 10110 ($1918 + 2 \times 4096$). If port 10110 is not available, the search goes to port 14206 ($1918 + 3 \times 4096$).

The agent is assigned to the first available port encountered in the search. The process fails to start if the search reaches the highest port without a successful binding (port 14206 in this example).

- The SKIP:N parameter specifies the number of ports to skip when starting the search for an available port using the port assignment algorithm. Use the SKIP parameter for components that do not need access across a firewall.

For example, if the well-known port is 1918, SKIP:2 specifies to start the search at port 10110 ($1918 + 2 \times 4096$), skipping ports 1918 and 6014 ($1918 + 1 \times 4096$). The algorithm continues searching until it finds an available port.

- The USE parameter enables or disables a protocol. To disable a protocol, specify use:n. To enable a protocol, specify use:y. This parameter has no default.

Note: Tivoli Monitoring agents allocate ports 1920 and 3661 as HTTP and HTTPS listener ports.

In environments behind firewalls it is critical that the Warehouse Proxy Agent listens on the same port each time it is started. To ensure a consistent IP.PIPE listening port, regardless of the sequence in which other monitoring agents on the system are started, the Warehouse Proxy Agent by default uses the high-numbered port 63358. This is accomplished by leveraging the COUNT and SKIP parameters. The Warehouse Proxy Agent internally sets COUNT:1 and SKIP:15 in order to obtain 63358 as its listening port ($63358 = 1918 + (15 \times 4096)$). You do not need to add these parameters when configuring the Warehouse Proxy Agent. All other monitoring agents that archive their historical data should then configure 63358 as the well-known listening port of the Warehouse Proxy Agent.

Note: You are not required to use port 63358 for the Warehouse Proxy Agent, but this is the most commonly used port. If you decide to use a different listening port, then you will need to add your own custom COUNT and SKIP values when configuring the agent.

Example

The example in Table 70 shows the coding to use on a system that contains the components shown:

Table 70. Using *COUNT* and *SKIP* variables to assign port numbers

Component	Coding
Tivoli Enterprise Monitoring Server	The monitoring server uses port 1918.
Warehouse Proxy Agent Requires firewall access	KDE_TRANSPORT=IP.PIPE COUNT:1 SKIP:15 This coding reserves port 63358 (1918 + 15*4096) for the Warehouse Proxy Agent.
Windows OS agent Does not require firewall access	KDE_TRANSPORT=IP.PIPE SKIP:2 This coding reserves port 10110 (1918 + 2*4096) for the Windows OS agent. If the agent fails to open port 10110, it tries port 14206 and so on, until it finds an available port or exhausts all possibilities.

Adding the KDE_TRANSPORT environment variable

The KDE_TRANSPORT environment variable must be added to the appropriate file (*ENV file on Windows).

Before you begin

The KDE_TRANSPORT variable supersedes and overrides the KDC_FAMILIES variable. If a KDC_FAMILIES variable exists in the file, merge the KDC_FAMILIES settings with the KDE_TRANSPORT settings. Copy the KDC_FAMILIES settings that you want to keep to the new KDE_TRANSPORT variable.

About this task

Use the following procedure to add the KDE_TRANSPORT environment variable to the appropriate file:

Add the KDE_TRANSPORT environment variable to the ENV file for the component.

For example, the ENV file for the Windows OS agent is named KNTENV, where NT is the product code for the Windows OS agent. For a list of product codes, see Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883.

Edit the ENV file:

Procedure

1. In the Manage Tivoli Enterprise Monitoring Services window, right-click the component that you want to change, point to **Advanced**, and select **Edit ENV File**.
2. Add a line similar to the following example:

```
KDE_TRANSPORT=IP.PIPE PORT:1918 COUNT:N SKIP:N  
IP use:n SNA use:n IP.SPIPE use:n
```

where *N* is the number of ports to reserve (COUNT:*N*) or the number of ports to skip (SKIP:*N*). This example uses the IP.PIPE protocol. It also applies to IP.SPIPE.
3. Search the file for a KDC_FAMILIES environment variable. If a KDC_FAMILIES variable exists in the file, merge its settings with the new KDE_TRANSPORT variable. The KDE_TRANSPORT variable supersedes and overrides the KDC_FAMILIES variable.
4. Save the file.

5. Recycle the component. (Right-click the component and select **Recycle**.)

Configuring the heartbeat interval

IBM Tivoli Monitoring uses a heartbeat mechanism to monitor the status of remote monitoring servers and monitoring agents.

The different monitoring components in the monitoring architecture form a hierarchy (shown in Figure 89) across which the heartbeat information is propagated.

The hub monitoring server maintains status for all monitoring agents. Remote monitoring servers offload processing from the hub monitoring server by receiving and processing heartbeat requests from monitoring agents, and communicating only status changes to the hub monitoring server.

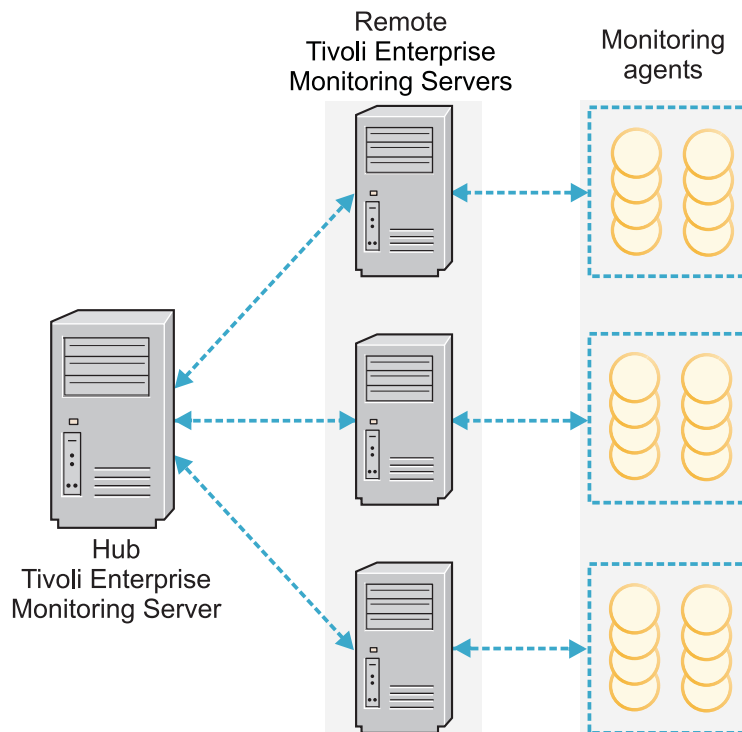


Figure 89. Hierarchy for the heartbeat interval

At the highest level, the hub monitoring server receives heartbeat requests from remote monitoring servers and from any monitoring agents that are configured to access the hub monitoring server directly (rather than through a remote monitoring server). The *default* heartbeat interval used by remote monitoring servers to communicate their status to the hub monitoring server is 3 minutes. The default heartbeat interval of 3 minutes for monitoring servers is suitable for most environments, and should not need to be changed. If you decide to modify this value, carefully monitor the system behavior before and after making the change.

At the next level, remote monitoring servers receive heartbeat requests from monitoring agents that are configured to access them. The *default* heartbeat interval used by monitoring agents to communicate their status to the monitoring server is 10 minutes.

You can specify the heartbeat interval for a node (either a remote monitoring server or a remote monitoring agent) by setting the `CTIRA_HEARTBEAT` environment variable. For example, specifying `CTIRA_HEARTBEAT=5` sets the heartbeat interval to 5 minutes. The *minimum* heartbeat interval that can be configured is 1 minute.

- For monitoring servers on Windows computers, you can set this variable by adding the entry to the KBBENV file. You can access this file from the Manage Tivoli Enterprise Monitoring Services utility by right-clicking **Windows OS Monitoring Agent** and clicking **Advanced -> Edit ENV File**. Note that you must stop and restart the monitoring server for the changes to the KBBENV file to take effect.
- For monitoring servers on Linux and UNIX computers, you can set the **CTIRA_HEARTBEAT** variable by adding the entry to the monitoring server configuration file. The name of the monitoring server configuration file is of the form *hostname_ms_temsname.config*. For example, a remote monitoring server named REMOTE_PPERF06 running on host pperf06 has a configuration filename of *ppperf06_ms_REMOTE_PPERF06.config*. Note that you must stop and restart the monitoring server for the configuration changes to take effect.
- For remote monitoring servers, you can set this variable by adding an entry to the KBBENV file. You can access this file from Manage Tivoli Enterprise Monitoring Services by right-clicking **Tivoli Enterprise Monitoring Server** and clicking **Advanced → Edit ENV File**. You must stop and restart the monitoring server for changes to the KBBENV file to take effect.
- For Windows OS agents, you can set this variable by adding the entry to the KNTENV file. You can access this file from Manage Tivoli Enterprise Monitoring Services by right-clicking **Windows OS Monitoring Agent** and clicking **Advanced → Edit ENV File**. You must stop and restart the monitoring agent for the changes to the KNTENV file to take effect.
- For agents on Linux and UNIX computers, you can set the **CTIRA_HEARTBEAT** variable by adding an entry to the agent .ini file (for example, lz.ini, ux.ini, ua.ini). When the agent is stopped and restarted, the agent configuration file is recreated using settings in the .ini file.

When a monitoring agent becomes active and sends an initial heartbeat request to the monitoring server, it communicates the desired heartbeat interval for the agent in the request. The monitoring server stores the time the heartbeat request was received and sets the expected time for the next heartbeat request based on the agent heartbeat interval. If no heartbeat interval was set at the agent, the default value is used.

Changes to offline status typically require two missed heartbeat requests for the status to change. Offline status is indicated by the node being disabled in the portal client's Navigator View. If the heartbeat interval is set to 10 minutes, an offline status change would be expected to take between 10 and 20 minutes before it is reflected on the portal client's Navigator View.

Attention: Lower heartbeat intervals increase CPU utilization on the monitoring servers processing the heartbeat requests. CPU utilization is also affected by the number of agents being monitored. A low heartbeat interval and a high number of monitored agents could cause the CPU utilization on the monitoring server to increase to the point that performance related problems occur. If you reduce the heartbeat interval, you must monitor the resource usage on your servers. A heartbeat interval lower than 3 minutes is not supported.

Restarting the Tivoli Enterprise Portal Server after reconfiguration

The Tivoli Enterprise Portal Server must be stopped before it can be reconfigured.

- If the portal server is running and you ask to reconfigure it, a warning is first displayed informing you that the server must be stopped and asking if you want to continue.
- If the portal server is running and you allow it to be stopped for reconfiguration, you are asked when the configuration process ends if you want the server restarted.

Verifying the dashboard data provider is enabled

Check the Tivoli Enterprise Portal Server configuration to verify the dashboard data provider is enabled by using Manage Tivoli Enterprise Monitoring Services or the command-line.

Procedure

- Using Manage Tivoli Enterprise Monitoring Services
 1. Start Manage Tivoli Enterprise Monitoring Services on the computer where the portal server is installed:
 - **Windows** Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Enterprise Monitoring Services**.
 - **Linux** **UNIX** Where *install_dir* is the IBM Tivoli Monitoring installation directory, change to the *install_dir/bin* directory and run `./itmcmd manage [-h install_dir]`.
 2. Right-click **Tivoli Enterprise Portal Server**:
 - **Windows** Click **Reconfigure**, and click **OK** to accept the existing configuration and go to the second TEP Server Configuration window.
 - **Linux** **UNIX** Click **Configure**. The Common Event Console Configuration window is displayed. Click **OK** to accept the current values. On the Configure Tivoli Enterprise Portal window, select the **Dashboard data provider** tab.
 3. In the dashboard data provider area of the configuration window, verify the **Enable the dashboard data provider** check box is selected. If it is not selected, then select it.

In this same area of the configuration window, you can specify a domain override value and enable authorization policies.

 - a. The domain override value is optional and is only valid when the dashboard data provider is enabled. It changes the default dashboard data provider ID and domain name for authorization policies from `itm.<hub monitoring server name>` to `itm.<domain override value>`. The value cannot exceed 124 characters. Best practice is to configure a domain override value for these scenarios:
 - The Hot Standby high availability feature is being used for the hub monitoring server. By configuring a domain override value, the dashboard data provider ID and domain name will not change when the portal server is configured to connect to the new acting hub monitoring server. If you do not configure a domain override value in this scenario, you must reconfigure the connection between the IBM Dashboard Application Services Hub and the dashboard data provider and update any domain specific authorization policies when the portal server is configured to connect to the new acting hub monitoring server.
 - You have multiple hub monitoring servers and you are setting up a dashboard environment with single sign-on with authorization policies, and you want to create some domain specific authorization policies. You must specify a domain override value for this scenario if you want to use a more user-friendly domain name in your authorization policies other than the default value of `itm.<hub monitoring server name>`.
 - b. The **Enable authorization policies** option is selected if you want to use authorization policies to control which managed systems and managed system groups a user can access in monitoring dashboards. Only enable authorization policies if you are setting up a dashboard environment with single sign-on, you plan to use authorization policies to control access to monitoring dashboards, and your administrators have already created the initial set of policies for dashboard user access.

Note: If you plan to use authorization policies and they have not been created yet, leave this box unchecked. After the authorization policies are ready, reconfigure the portal server and enable authorization policies.
 4. If you made changes to the portal server configuration, click **OK**. Otherwise, click **Cancel** to exit.
 5. If you clicked **OK** and are prompted to reconfigure the warehouse connection information, answer **No**.
 6. On Windows, after some processing of the configuration settings, the Common Event Console Configuration window is displayed. Sometimes this window does not open in the foreground and

is hidden by other windows. If processing seems to be taking longer than expected, minimize other windows and look for the configuration window. When the Common Event Console Configuration window is displayed, click **OK**.

7. If you made configuration changes, ensure the portal server is restarted.

- Using the command-line

If the Tivoli Enterprise Portal Server is on Linux or UNIX, you can check the portal server configuration from the command-line and enable the dashboard data provider if it is not already enabled.

1. Log on to the computer where the Tivoli Enterprise Portal Server is installed.
2. At the command-line, change to the *install_dir/bin* directory, where *install_dir* is the directory where you installed the product.

3. Run the following command to view the Tivoli Enterprise Portal Server configuration settings:
./itmcmd config -g -A cq.

4. Check if the `KD8_ITM_REST_SERVICE_ENABLE` variable is set to Y.

If is not present, or is set to N, then the dashboard data provider is not enabled, so proceed with the remaining steps in this procedure.

If it is set to Y, you have successfully verified the dashboard data provider is enabled and can skip the remaining steps.

5. Run the following command to configure the Tivoli Enterprise Portal Server: **./itmcmd config -A cq.**

The message Agent configuration started is displayed, followed by a prompt: Tivoli Enterprise Portal Server will be stopped during configuration. Do you want to continue? [1=Yes, 2=No] (Default is: 2).

6. Enter 1. The following prompt is displayed: Edit "Common event console for IBM Tivoli Monitoring" settings? [1=Yes, 2=No] (default is: 1).

7. Enter 2. The following prompt is displayed: Will this agent connect to a TEMS? [1=YES, 2=NO] (Default is: 1).

8. Accept the default values for this prompt and the prompts that follow it until you see the prompt about configuring the dashboard data provider. If it is not enabled, select a value of 1 to enable it.

9. Next you are asked if you want to specify a domain override value.

- a. The domain override value is optional and is only valid when the dashboard data provider is enabled. It changes the default dashboard data provider ID and domain name for authorization policies from `itm.<hub monitoring server name>` to `itm.<domain override value>`. The value cannot exceed 124 characters. Best practice is to configure a domain override value for these scenarios:

- The Hot Standby high availability feature is being used for the hub monitoring server. By configuring a domain override value, the dashboard data provider ID and domain name will not change when the portal server is configured to connect to the new acting hub monitoring server. If you do not configure a domain override value in this scenario, you must reconfigure the connection between the IBM Dashboard Application Services Hub and the dashboard data provider and update any domain specific authorization policies when the portal server is configured to connect to the new acting hub monitoring server.
- You have multiple hub monitoring servers and you are setting up a dashboard environment with single sign-on with authorization policies, and you want to create some domain specific authorization policies. You must specify a domain override value for this scenario if you want to use a more user-friendly domain name in your authorization policies other than the default value of `itm.<hub monitoring server name>`.

10. If the dashboard data provider is enabled, you are prompted whether you want to enable authorization policies. Only enable authorization policies if you are setting up a dashboard

environment with single sign-on, you plan to use authorization policies to control access to monitoring dashboards, and your administrators have already created the initial set of policies for dashboard user access.

Note: If you plan to use authorization policies and they have not been created yet, leave this box unchecked. After the authorization policies are ready, reconfigure the portal server and enable authorization policies.

11. After the installer has completed the configuration, the following message is displayed: Agent configuration completed and you are asked if you want to restart the portal server. Select 1 to restart it.

Results

You have successfully verified that the dashboard data provider is enabled

Switching to a different Tivoli Enterprise Portal Server database

You can change to a different RDBMS after you have installed the Tivoli Enterprise Portal Server.

Procedure

1. Bring up the Manage Tivoli Enterprise Monitoring Services main screen, and right-click the Tivoli Enterprise Portal Server entry. See Figure 90.

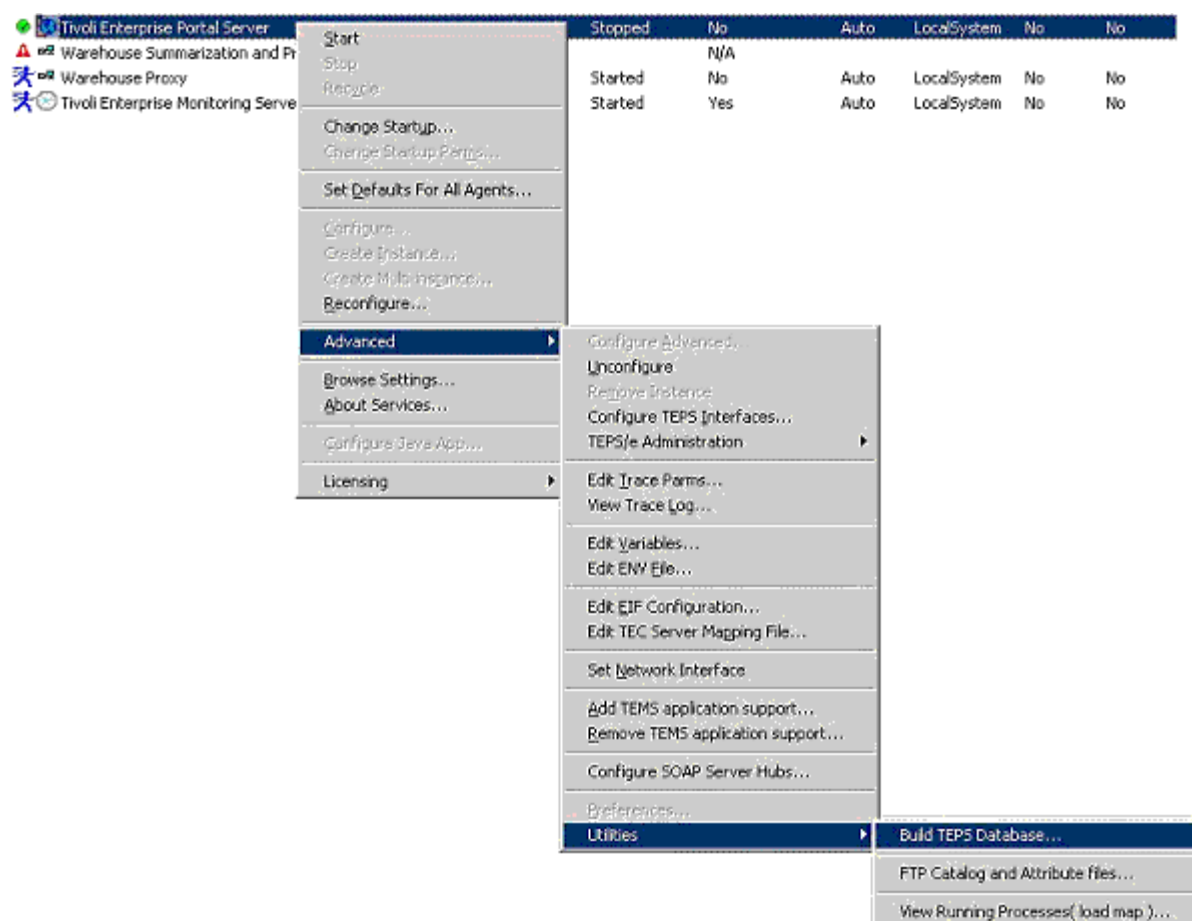


Figure 90. Manage Tivoli Enterprise Monitoring Services Advanced Utilities window

2. From the pop-up menu, select Advanced → Utilities → Build TEPS Database.
 - If the only RDBMSes installed on this computer are DB2 Database for Linux, UNIX, and Windows and the portal server's embedded Derby database, select the appropriate database manager, as shown in Figure 91.



Figure 91. The Manage Tivoli Enterprise Monitoring Services select the new portal server database window. The only available database managers are DB2 for Linux, UNIX, and Windows and Derby.

- If this computer is running both DB2 Database for Linux, UNIX, and Windows, Microsoft SQL Server, and embedded Derby, select the appropriate database manager, as shown in Figure 92.



Figure 92. The Manage Tivoli Enterprise Monitoring Services select the new portal server database window. The available database managers are DB2 for Linux, UNIX, and Windows, SQL Server, and Derby.

Note: Changing to a different database does not migrate the portal server data stored in it.

3. Continue with the database configuration, as explained in step 12 on page 173 under “Installation procedure” on page 172.

Silent configuration of the Performance Analyzer

The silent method of configuration is useful for advanced users who want to supply configuration information once through the silent configuration file for the Performance Analyzer.

About this task

With silent configuration, you can automate the configuration steps and avoid manual configuration of the Performance Analyzer.

Procedure

1. Locate the default silent configuration file for the Performance Analyzer.
 - **Windows** One sample `silent_server.txt` file is used by all agents for silent configuration. The sample file is provided on the product installation media.

- **Linux** **UNIX** Each agent has a separate configuration file. For Performance Analyzer, the default configuration file is `pa_silent_config.txt`. For example: `<ITM_HOME>/samples/pa_silent_config.txt`.
- 2. Review the default configuration parameters in the silent configuration file. If required, update the values according to your requirements.
- 3. Save the file with a different file name.
- 4. Run the Performance Analyzer silent configuration command:
 - **Windows**
 - For silent configuration at Performance Analyzer installation, go to the IBM Tivoli Monitoring installation directory and run the command:


```
start /wait setup /z"/sf<Complete path to silent_server.txt>" /s /f2"<complete path to log file>"
```
 - For silent configuration after Performance Analyzer installation, create the silent configuration file by copying Performance Analyzer-specific parameters and monitoring server connection-related parameters from the `silent_server.txt` file. Use the new silent configuration file in the command:


```
<ITM_HOME>\Install\ITM\kinconfg.exe -ak<pc> -n<Complete path_to_silent_file>
```

Where `<pc>` = `pa`
 - **Linux** **UNIX** Run the following command (from the `bin` directory of Tivoli Monitoring if that is not added to `PATH`):


```
itmcmd config -A -h <ITM_HOME> -p <Complete path to silent_config_pa.txt> pa
```

Configuring historical data collection for the Performance Analyzer warehouse agent

You must configure the historical data collection for Tivoli Performance Analyzer to enable trend calculation visualization in the workspaces and reports.

About this task

Complete the following steps to configure historical data collection for the Performance Analyzer warehouse agent:

Procedure

1. Start Tivoli Enterprise Portal.
2. In the portal toolbar, select **Edit > History Configuration**. The **History Collection Configuration** window is displayed.
3. To obtain the calculation data, configure the historical data collection for each of your operating system agents, and for your domains by referring to the attribute groups that must be configured. Table 71 on page 433 shows products and their associated attribute groups.

Table 71. Attribute groups that must be configured

Product		Attribute groups
Operating systems	Windows OS	<ul style="list-style-type: none"> • Network Interface (Superseded) • Logical Disk • Memory (Superseded) • Processor • System
	Linux	<ul style="list-style-type: none"> • Linux CPU Averages (Superseded) • Linux Disk (Superseded) • Linux Network (Superseded) • Linux System Statistics (Superseded) • Linux VM Stats (Superseded)
	UNIX OS	<ul style="list-style-type: none"> • Disk • Network • System
DB2		<ul style="list-style-type: none"> • Database (Superseded) • System Overview (Superseded) • Tablespace (Superseded)
Oracle		<ul style="list-style-type: none"> • Oracle Cache Totals • Oracle Database • Oracle Library Cache Usage • Oracle Server • Oracle Session Summary • Oracle Statistics Summary • Oracle Tablespaces
ITCAM for RT	Client response time	<ul style="list-style-type: none"> • CRT Application Status • CRT Server Status • CRT SubTransaction Status • CRT Transaction Status
	Robotic response time	<ul style="list-style-type: none"> • RRT Application Status • RRT Robotic Playback Status • RRT SubTransaction Status • RRT Transaction Status
	Web response time	<ul style="list-style-type: none"> • WRT Application Status • WRT Client Status • WRT Server Status • WRT SSL Alert Current® Status • WRT SubTransaction Status • WRT TCP Status • WRT Transaction Status

Table 71. Attribute groups that must be configured (continued)

Product		Attribute groups
System	VIOS Premium	<ul style="list-style-type: none"> • KVA LOGICAL PARTITION • KVA DISKS • KVA NETWORK ADAPTERS RATES
	AIX Premium	<ul style="list-style-type: none"> • KPX LOGICAL PARTITION
VMware		<ul style="list-style-type: none"> • KVM SERVER • KVM SERVER NETWORK • KVM VM CPU • KVM VM PARTITION • KVM VM NETWORK • KVM VM MEMORY

For each agent, do the following steps:

- a. In the navigation pane, right-click the product whose agent you want to configure, and select **Create new collection settings**. The **Create New Collection Settings** window opens.
 - b. Select an attribute group for that agent, as listed in Table 71 on page 433, and click **OK**.
 - c. In the **Basic** tab, specify the following settings:
 - In the **Collection Interval** field, set the polling interval, for example to **1 hour**.
 - In the **Warehouse Interval** field, set the frequency at which data is written to the database.
 - d. In the **Distribution** tab, specify for which nodes the data should be collected and click **OK**.
 - e. Select the product again, and set its **Summarization** interval to at least **Hourly** and **Daily**. You can also select additional summarization intervals.
 - f. Set the **Pruning** interval to what is most appropriate for your system, and click **OK**.
 - g. Repeat the preceding steps for each agent support agent in your environment. Use collection, summarization, and pruning values that are appropriate to your system. You have now completed configuring the support agents for your environment.
4. Next, configure the historical data collection for Performance Analyzer Warehouse Agent to enable trend calculation visualization on the workspaces and reports:
 - a. In the navigation pane, click Performance Analyzer Warehouse Agent.
 - b. Select the first attribute group for the agent.
 - c. In the **Basic** tab, specify the following settings:
 - In the **Collection Interval** field, set the polling interval to **1 hour** for System Health and Disk Health attribute groups and to **1 day** for other attribute groups.
 - In the **Warehouse Interval** field, set the interval to **1 day**.
 - d. In the **Distribution** tab, start collection for the **AFF_PERF_ANALYZER_WHSE_AGENT** group.
 - e. Repeat steps c and d for all attribute groups.
 - f. Select the Performance Analyzer Warehouse Agent application and set **Pruning** for detailed data for all attribute groups to **1 day**.
 - g. Finally, click **Close** when you have completed configuring the historical data collection for the agents.

Results

1. The out-of-the-box OS domain uses 32-bit attribute groups only – the metrics used in this domain are expressed in units that never exceed maximum value available for 32-bit numbers. There is no additional value in switching OS domain tasks, workspaces, and reports to use 64-bit attribute groups.

2. For step 4 on page 434: instead of using History Collection Configuration dialog in Tivoli Enterprise Portal to set up warehousing and pruning of Tivoli Performance Analyzer attribute groups, it is possible to achieve the same with the **tacmd** command. Sample scripts to automate this task can be found on the installation media under the **scripts/<domain_name>** directory. For example, the OS domain scripts are:

```
scripts/os/histcoll_os.sh  
or  
scripts/os/histcoll_os.bat
```

Tivoli Monitoring protocol usage and protocol modifiers

IBM Tivoli Monitoring basic services communications are defined by the KDE_TRANSPORT environmental variable.

In earlier releases KDC_FAMILIES was used but KDE_TRANSPORT is the successor. The two variables are processed identically, however many agent installers are only aware of the earlier variable. You should examine the agent installation and adopt whatever is used.

IBM Tivoli Monitoring uses other communication protocols such as the following:

- Tivoli Enterprise Portal Client to Tivoli Enterprise Portal Server uses CORBA IIOP after startup.
- Tivoli Enterprise Portal Server, Warehouse Proxy Agent, and Summarization and Pruning Agent can use ODBC or JDBC to communicate with the warehouse database.
- Most IBM Tivoli Monitoring to IBM Tivoli Monitoring communications use the TCP/IP protocol.

Modifiers to the protocols are of the form attribute:value. If the modifier occurs first, then it is global in effect. If the modifier occurs after a protocol and before the next protocol, then the effect is only on the protocol that precedes it. The protocol names and modifiers are not case sensitive although they are presented in uppercase in the following paragraphs.

KDE_TRANSPORT Structure

KDE_TRANSPORT Structure is a string that lists protocols and modifiers. All protocols are assumed as present. A protocol is activated only if an interface is available for use.

The USE modifier activates a specific protocol if Y is specified, otherwise it deactivates it. The scanning starts with an implicit global USE:Y, meaning that all protocols are assumed to be activated by default. For example:

```
IP.PIPE PORT:1918 USE:Y IP.SPIPE USE:N
```

This modifier means that IP.PIPE will be available but not IP.SPIPE. In addition, all the unnamed protocols listed later are activated.

KDE_TRANSPORT Transmission Control Protocol

Transmission Control Protocol (TCP) is a connection oriented protocol. Connection is made through a port number such as 0-65535.

The connection continues until the application takes it down. Here are the protocol names:

```
IP.PIPE - tcp  
IP.SPIPE - secure tcp  
IP6.PIPE - ipv6 tcp  
IP6.SPIPE - ipv6 secure tcp
```

The secure protocols are implemented with the Global Secure Toolkit (GSKIT) component.

The TCP protocol modifiers are listed in the sections that follow.

PORT

PORT defaults to 1918 for TCP and 3660 for secure TCP.

These numbers are registered with the Internet Assigned Numbers Authority, for more information go to <http://www.iana.org/assignments/port-numbers>.

PORT defines the base port number. For a Tivoli Enterprise Monitoring Server, the base number is the listening port and the port that agents connect to. For agents, IBM Tivoli Monitoring processing attempts to open a listening port at number $\text{base} + N * 4096$, where N is a number from 1 to 15. If one is already in use Tivoli Monitoring attempts to open a listening port at the next higher iteration. If no Tivoli Enterprise Monitoring Server is present the base port is reserved, in the event a Tivoli Enterprise Monitoring Server will be started later.

An agent uses the listening port for two main purposes:

- Tivoli Enterprise Monitoring Server requesting real-time data from the agent,
- An agent receiving notifications from the Tivoli Enterprise Monitoring Server, such as awareness of a Warehouse Proxy Agent re-registration at a new IP address or port number.

Example: `IP.PIPE PORT:1918 USE:Y`

Note: In this definition mode there are a maximum of 15 agents on a server. If you have a scenario with more than 15 agents on a server, see the **EPHEMERAL** modifier section.

SKIP and COUNT

SKIP and COUNT modifiers are used to control the port search algorithm. Default search is for $\text{baseport} + N * 4096$ where N is a number from 1 to 15. The SKIP modifier forces it to start with N equal to the SKIP value in the above calculation. The COUNT modifier controls the number of attempts that are made.

The following modifier is commonly used for the Warehouse Proxy Agent:

`IP.PIPE PORT:1918 SKIP:15 COUNT:1 USE:Y`

The only port that will be checked with this modifier is $1918 + 15 * 4096$ or 63358. This means that the Warehouse Proxy Agent will have a fixed IP address even when a Tivoli Enterprise Monitoring Server and other agents are starting up. Having a fixed port number is required if firewall rules are in place.

EPHEMERAL

This section explains the three EPHEMERAL modifier values.

This modifier has three different values:

- A value of Y means that the connection to the Tivoli Enterprise Monitoring Server listening port is used for all communications. The agent does not require a separate listening port. This means fewer ports are used, which can be important when firewall rules are in place. It is also a way to avoid the 15 agent limitation. There is a disadvantage in the case of historical data collection. The historical data must be either stored at the Tivoli Enterprise Monitoring Server or, if the historical data is being stored at the agent, a Warehouse Proxy Agent must be running on the same server as the Tivoli Enterprise Monitoring Server that the agent is reporting to.
- A value of OUTBOUND means the same as the value Y above.
- A value of INBOUND can be used at a Tivoli Enterprise Monitoring Server. It means that every agent connecting to the Tivoli Enterprise Monitoring Server is configured to ephemeral mode.

POOL

IBM Tivoli Monitoring processes use opaque ports (also sometimes referred to as ephemeral ports) for communication with the server.

These ports are not visible on the outside network. The range of possible opaque port numbers that are used can be controlled with the POOL option. See the following examples:

```
IP.PIPE POOL:50900-51923
IP.UDP POOL:01000-01023 POOL:01024-02048
```

Note: Each pool modifier can specify a maximum of 1024 pool ports, but you can have more than one such specification.

KDE_TRANSPORT User Datagram protocol

Typically this is not the best choice of protocol for IBM Tivoli Monitoring. Its advantage is a somewhat lower storage requirement.

The disadvantages are as follows:

- Less reliability since applications are responsible for error recovery.
- Higher CPU resources are required.
- It cannot be used where firewall rules are in place.

The protocol names are as follows:

```
IP.UDP - User Datagram Protocol
IP - Synonym for IP.UDP
IP6.UDP - IP V6 version of IP.UDP
```

The POOL and PORT modifiers can be used for this scenario. The other modifiers are connection oriented.

KDE_TRANSPORT Hypertext Transfer Protocol

Each IBM Tivoli Monitoring process has an internal web server. The KDE_TRANSPORT Hypertext Transfer Protocols define what protocols are used to access the internal web server. The web server provides access to the IBM Tivoli Monitoring Service Console, the Tivoli Enterprise Portal client, the SOAP server tryout page, and the Agent Service Index pages (with Tivoli Monitoring V6.2.2 or higher).

The protocol names are as follows:

```
ip.tcp.http - http communications
ip.ssl.https - secure http communications
ip6.tcp.http - ipv6 http communications
ip6.ssl.https - ipv6 secure http communications
```

The protocol modifiers are as follows:

- HTTP_SERVER: Defaults to Y. If set to N then the internal web server is not started.
- HTTP_CONSOLE: Defaults to Y. If set to N then the IBM Tivoli Monitoring service console is not started.
- HTTP: Defaults to 1920. If set to 0, access to non-secure internal web server is disabled.
- HTTPS: Defaults to 3661. If set to 0, access to secure internal web server is disabled.

Note: To meet the FIPS 140-2 requirement for secure communication, you should consider setting HTTP:0 and only allowing internal web server access with HTTPS.

- POOL: The HTTP protocol also uses temporary ports and the usage is controlled by separate pool control settings. These are not protocols, but are required for the POOL setting:

```
ip.tcp - pool control for the ip.tcp.http protocol
ip.ssl - pool control for the ip.ssl.https protocol
ip6.tcp - pool control for the ip6.tcp.http protocol
ip6.ssl - pool control for the ip6.ssl.https protocol
```

Here is an example of usage: IP.TCP.HTTP USE:Y IP.TCP POOL:20000-20031. You do not need to specify a protocol that is assumed enabled already, so you could simply add: IP.TCP POOL:20000-20031. To completely control pool usage, you must set the POOL value for all enabled protocols.

How to change KDE_TRANSPORT settings

If you decide to change KDE_TRANSPORT you must change the settings on all IBM Tivoli Monitoring tasks, when appropriate.

For example, if you decide to turn off the IBM Tivoli Monitoring internal web server using the HTTP_SERVER:N modifier, you must make the same change for all IBM Tivoli Monitoring tasks.

Mass change to Windows and Linux/UNIX agents

For details on how to make a mass change to Windows and Linux/UNIX agents, see <http://www-01.ibm.com/support/docview.wss?uid=swg21441836>.

Linux/UNIX – Tivoli Enterprise Monitoring Server

This environment is similar to agents. The difference is the config file generation is only performed when you run the command:

```
itmcmd config -S -t <temsname>
```

Edit the ms.ini in the same way as the agent case above, and then reconfigure the Tivoli Enterprise Monitoring Server using the itmcmd config command.

i/5

The environment variable is manually changed in:

```
QAUTOTMP/KMSPARM(KBBENV)
```

z/OS

In z/OS these values are kept in the RKANPARU(KDSENV) member for Tivoli Enterprise Monitoring Server and in RKANPARU(KppENV) for agents. Configuration changes here are manual.

Interactions with other environment variables

KDEB_INTERFACELIST and KDEB_INTERFACELIST_IPV6

The dash “-” option is used alone. These environment variables do not scan any of the related interfaces. In such a scenario a protocol might go unused even though it is specified in KDE_TRANSPORT. You can also eliminate all the interfaces by name to achieve the same result.

Chapter 14. Additional Linux and UNIX configuration steps

This chapter provides information about additional configuration that might be required for Linux and UNIX systems.

See the following sections for related information:

- “Disabling `fsync()` calls”
- “Configuring permissions for a monitoring server on a non-NIS Solaris computer”
- “Linux requirements for the localhost host name” on page 440
- “Setting `ulimit` values for the Warehouse Proxy Agent” on page 440

Disabling `fsync()` calls

The `KGLCB_FSYNC_ENABLED` parameter can be used to specify whether the `fsync()` system call should be invoked after writes to the filesystem. This configuration variable may be set in the standard configuration file for the monitoring server.

For maximum reliability, by default `fsync()` is called. The `fsync()` system call flushes the filesystem's dirty pages to disk and protects against loss of data in the event of an operating system crash, hardware crash, or power failure. However, the call can have a significant negative effect on performance, because in many cases it defeats the caching mechanisms of the platform filesystem. On many UNIX platforms, the operating system itself syncs the entire filesystem on a regular basis. For example, by default the `syncd` daemon that runs on AIX syncs the filesystem every 60 seconds, which limits the benefit of `fsync()` calls by application programs to protecting against database corruption in the most recent 60-second window.

If the following line is added to the monitoring server configuration file, `fsync()` calls are omitted:

```
KGLCB_FSYNC_ENABLED='0'
```

Configuring permissions for a monitoring server on a non-NIS Solaris computer

If your monitoring server is installed on a non-NIS Solaris computer, you must set the permissions. You do not need to do this for monitoring servers running on AIX or Linux.

About this task

Set permissions for a monitoring server on a non-NIS Solaris system as follows:

Procedure

1. Go to the `/bin` directory where file `kdsvlunx` is located (*install_dir/arch/ms/bin*, where *arch* is the operating system on which the monitoring agent was installed).
2. Move the file to the root user ID if you have the root password; otherwise obtain the password from an administrator:

```
su root
chown root kdsvlunx
chmod u+s kdsvlunx
```
3. Return to your regular ID after you have moved the user ID to root.

Linux requirements for the localhost host name

KDH servicepoint manipulation and the kdh (http/https) port-sharing code requires the availability of localhost: the **localhost** host name must exist and be resolvable to an address on the local system.

Typically, **localhost** resolves to the loopback device with IP address 127.0.0.1. If **localhost** does not exist or cannot be resolved to an IP address, http-based servers (the SOAP server, the IBM Tivoli Monitoring Service Console, the Tivoli Enterprise Portal Server) will fail to initialize.

When **localhost** resolves to an address other than 127.0.0.1, a network interface must be locally available and configured to IP with that **localhost** network address.

Setting ulimit values for the Warehouse Proxy Agent

There is a limit on the number of open file descriptors that can be opened by the Warehouse Proxy process when it is running on UNIX and Linux systems. When a Tivoli Monitoring operating system monitoring agent connects to the Warehouse Proxy Agent, it uses a set of file descriptors for communication. When the number of operating system agent connections exceeds the file descriptor limit, the agent process consumes high amounts of CPU as it is unable to send data.

To correct this high CPU situation, the **ulimit** value must be set to a value higher than the maximum number of file descriptors that could be opened to the Warehouse Proxy Agent. The value that should be used is based upon the following conditions:

- The minimum number of file descriptors used by the Warehouse Proxy Agent process is $X + Y + 10$, where:

- | | |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| X | Is the number of agents warehousing to the Warehouse Proxy Agent |
| Y | Is the number of database connections between Warehouse Proxy Agent and the database server. (The default number of database connections is 10; you can change this number.) |
| 10 | Is the number of log and configuration files used by the Warehouse Proxy Agent. |

The value used for the file descriptor **ulimit** should be high enough so that the limit is not met. A simpler formula is $X + 1000$, where X is the number of agents warehousing.

After you determine the value for the file descriptor **ulimit**, modify the **ulimit** as appropriate for the operating system. See the system documentation for the command (usually **ulimit**) and procedures to make this change permanently across system restarts, or contact your UNIX or Linux System Administrator.

Chapter 15. Configuring IBM Tivoli Monitoring components for IPv6 communication

This chapter provides instructions for configuring various Tivoli Monitoring components for IP version 6 (IPv6) communication.

To enable IPv6 communication between a pair of endpoints, follow the instructions for configuring each endpoint. For example, if you want to enable IPv6 communication between a hub monitoring server on a Linux computer, and a remote monitoring server on Windows, follow the instructions for configuring the hub on the Linux computer and then the remote server on the Windows computer. For ease of reference, instructions are provided for each platform separately.

Note: In order to use IPv6 with the IBM Tivoli Monitoring components, you must ensure your operating system is set up and configured for IPv6.

For details on supported combinations of IPv6 with IPv4 across the various IBM Tivoli Monitoring components, see “Choose between IPv6 and IPv4” on page 127.

By default, Tivoli Monitoring uses the well known ports for communication between components. If the default ports that are used by Tivoli Monitoring components in the environment are changed, the change must be made consistently to the KDC_FAMILIES/KDE_TRANSPORT variables for all components, specifying the new port parameter for various protocols. For instance, if the port 11111 is to be used instead of the default 1918, and 7663 instead of the default 3660, the KDC_FAMILIES variable is as follows:

```
KDC_FAMILIES=ip6.pipe port:11111 ip6 port:11111 ip6.spipe port:7663 ....
```

If IPv4 is also used, then the specification for the IPv4 protocol is also included. The syntax is platform-dependent.

Configuring AIX, Solaris, HP-UX, and Linux monitoring components

The instructions in this section apply to configuring monitoring components on the AIX, Solaris, HP-UX, and Linux platforms.

\$CANDLEHOME refers to the directory in which the Tivoli monitoring product is installed. In the following example, if a numeric IPv6 address is provided instead of a host name, the address must be enclosed in parentheses:

```
export CT_CMSLIST='ip6.pipe:(2002:930:9B04:305:9:48:133:98)'
```

Configuring the hub monitoring server

You must first configure hub and remote monitoring servers using the platform-specific installation and configuration tools, which includes selecting a protocol and providing information about the standby hub.

If IPv6 is the preferred protocol for the hub, then the glb_site.txt file must be edited, providing the desired IPv6 family as the protocol. For example:

```
ip6.pipe:proton
```

The glb_site file is found in the tables directory for the hub monitoring server, for example, \$CANDLEHOME/tables/\$TEMSNAME.

On all platforms, if a port other than the default is used, then the KDC_FAMILIES variable must be added or updated in the configuration file for the hub in the \$CANDLEHOME/config directory:

```
Export KDC_FAMILIES='ip6.pipe port:11111 ip6.port:11111 ip6.spipe port:7663 ....'
```

On AIX only, the IPv6 protocol must be explicitly provided in the KDC_FAMILIES variable, before any IPv4 protocols. This variable is found in the configuration file for the hub in the \$CANDLEHOME/config directory. For instance:

```
export KDC_FAMILIES='ip6.pipe use:y ip6.spipe use:y  
ip6.tcp use:y ip6.ssl use:y ip.pipe use:y ip.spipe use:y  
http use:y https use:y'
```

or simply:

```
export KDC_FAMILIES='+ip6 +ipv4'
```

However, if a port other than the default is being used, the port must be explicitly declared, as described at the beginning of this paragraph.

On Solaris only, the hub can be configured either for IPv4 or IPv6 but not both. For IPv4, no action is required. For IPv6, the following must be added to the configuration file for the hub in the \$CANDLEHOME/config directory:

```
export KDEB_INTERFACELIST='-'
```

On HP-UX on HP 9000 only, the hub can be configured either for IPv4 or IPv6 but not both. For IPv4, no action is required. For IPv6, the following must be added to the configuration file for the hub in the \$CANDLEHOME/config directory:

```
export KDEB_INTERFACELIST='-'
```

If the hub has been configured to use a standby, then the MHM:MIRROR_SITE variable must be edited if the preferred protocol for communication with the standby hub is IPv6. For example:

```
export MHM:MIRROR_SITE='ip6.pipe:atlas'
```

This variable is found in the KBBENV file for the monitoring server in the tables directory, for example, \$CANDLEHOME/tables/\$TEMSNAME.

If the server has been configured to provide access to other SOAP server hubs, then the kshxhubs.xml file must be updated to enable IPv6 communication with those hubs. This file can be found in the HTML directory for the monitoring server under the tables directory, for example, \$CANDLEHOME/tables/\$TEMSNAME/HTML.

Example:

```
<?xml version="1.0"?>  
<?xml-stylesheet type="text/xsl" href="hubdef.xsl"?>  
<ENTERPRISE>  
<HUB>  
<CMS_Name>  
ip6.pipe:saturn[1918]  
</CMS_Name>  
<Service_Name>newhub</Service_Name>  
<Alias>newhub</Alias>  
</HUB>  
</ENTERPRISE>
```

If the server has been configured to use LDAP for user authentication, the server uses available protocols to connect to the LDAP server. The LDAP server host can be designated by name or IP address. For example:

```
KGL_LDAP_HOST_NAME='saturn'  
KGL_LDAP_HOST_NAME='2002:930:9b04:305:9:48:157:7'
```

Configuring the Tivoli Enterprise Monitoring Automation Server

You must change the existing KDC_FAMILIES variable in the Tivoli Enterprise Monitoring Automation Server configuration file to match the added or updated KDC_FAMILIES variable in the hub monitoring server configuration file. In addition, you must also add the automation server connection protocol and port values to the KDC_FAMILIES variable for the automation server.

When adding or updating the KDC_FAMILIES variable in the hub monitoring server configuration file for IPv6 communication as described in “Configuring the hub monitoring server” on page 441, you must apply these same updates to the KDC_FAMILIES variable in the Tivoli Enterprise Monitoring Automation Server configuration file (as.ini) in the \$CANDLEHOME/config directory.

For example, if the automation server connection protocol and port values are HTTP and 10001, and the KDC_FAMILIES variable for the hub monitoring server was set to the following:

```
Export KDC_FAMILIES='ip6.pipe port:11111 ip6 port:11111 ip6.spipe port:7663'
```

Then edit the as.ini file in \$CANDLEHOME/config and change the existing KDC_FAMILIES variable from:

```
KDC_FAMILIES=$NETWORKPROTOCOL$
```

to

```
KDC_FAMILIES=ip6.pipe port:11111 ip6 port:11111 ip6.spipe port:7663 HTTP:10001
```

Note: The automation server will not start if there is an incorrect value for the protocol or port defined in either the KDC_FAMILIES or KDE_TRANSPORT environment variables. Verify that the settings support either HTTP or HTTPS protocol and a port value between 1 and 65535.

Configuring the remote monitoring server

You must first configure hub and remote monitoring servers using the platform-specific installation and configuration tools, which includes selecting a protocol and providing information about the standby hub.

The remote monitoring server uses the entries in the glb_site.txt file to look up the hub monitoring server. If there is a standby hub, then there are at least two entries in this file. This file must be edited if the protocol for communicating with one or both the hubs must be set to IPv6. For example:

```
ip6.pipe:proton  
ip.pipe:amoeba
```

On all platforms, if a port other than the default is used, then the KDC_FAMILIES variable must be added or updated in the configuration file for the hub the \$CANDLEHOME/config directory:

```
export KDC_FAMILIES='ip6.pipe port:11111 ip6 port:11111 ip6.spipe port:7663 ....'
```

On AIX only, the IPv6 protocol must be explicitly provided in the KDC_FAMILIES variable, before any IPv4 protocols. For instance:

```
export KDC_FAMILIES='ip6.pipe use:y ip6.spipe use:y  
ip6.tcp use:y ip6.ssl use:y ip.pipe use:y ip.spipe use:y  
http use:y https use:y'
```

or simply:

```
export KDC_FAMILIES='+ip6 +ip4'
```

However, if a port other than the default is being used the port must be explicitly declared, as previously described.

On Solaris only, the monitoring server can be configured either for IPv4 or IPv6 but not for both. For IPv4, no action is required. For IPv6, the following must be added to the configuration file for the monitoring server in the configuration directory:

```
export KDEB_INTERFACELIST='-'
```

On HP-UX on HP 9000 only, the monitoring server can be configured either for IPv4 or IPv6 but not for both. For IPv4, no action is required. For IPv6, the following must be added to the configuration file for the monitoring server in the configuration directory:

```
export KDEB_INTERFACELIST='-'
```

In this case, you must set the protocol to ip6.pipe in the glb_site file for this remote monitoring server.

Configuring the monitoring agents

Agents use the CT_CMSLIST variable to locate the monitoring server they are required to connect to.

This variable may contain one or more entries, separated by a semi-colon. For example:

```
CT_CMSLIST='ip6.pipe:proton;ip6.pipe:atlas'
```

If IPv6 is the preferred protocol for the agent to communicate with the designated monitoring server, then this variable must be updated. Before doing this, ensure that the monitoring server is configured to communicate using IPv6.

The procedure includes the following steps:

1. Open the xx.ini file in the \$CANDLEHOME/config directory for editing, where xx is the product code for the agent, such as lz.
2. Add the following line as the last line in the file (note that the line begins with a dot followed by a space):

```
. $CANDLEHOME$/config/xx.ipv6config
```
3. Repeat Steps 1 and 2 for the xx.config file in the same directory.
4. Create the xx.ipv6config file in the \$CANDLEHOME/config directory with the CT_CMSLIST specification, as in the following example:

```
export CT_CMSLIST='ip6.pipe:proton;ip6.pipe:atlas'
```

On all platforms, if a port other than the default is used, then the KDC_FAMILIES variable must be added in the xx.ipv6config file:

```
export KDC_FAMILIES='ip6.pipe port:11111 ip6 port:11111 ip6.spipe port:7663 ip.pipe port:11111 ....'
```

On AIX only, the IPv6 protocol must be explicitly provided in the KDC_FAMILIES variable, before any IPv4 protocols. For instance:

```
export KDC_FAMILIES='ip6.pipe use:y ip6.spipe use:y  
ip6.tcp use:y ip6.ssl use:y ip.pipe use:y ip.spipe use:y  
http use:y https use:y'
```

or simply:

```
export KDC_FAMILIES='+ip6 +ip4'
```

This definition must be provided in the xx.ipv6config file. However, if a port other than the default is being used, the port must be explicitly declared, as previously described.

On AIX, if the agent being configured is based on Tivoli Monitoring V6.1 but installed in a Tivoli Monitoring V6.3 environment, an additional step is required. For existing Tivoli Monitoring V6.1 agents, the preferred approach is to upgrade these agents to the Tivoli Monitoring V6.3 level. The upgrade can be done as follows:

- If there is a Tivoli Monitoring V6.1 Linux, HP-UX, Solaris, or AIX agent in the same environment, this agent can be upgraded to Tivoli Monitoring V6.3.
- If there is no Linux, HP-UX, Solaris, or AIX agent in the same environment, a Tivoli Monitoring V6.3 agent can be installed in this environment.

If the OS agent cannot be installed or upgraded in the agent environment for some reason, you should contact IBM support for assistance.

Check the LIBPATH variable in the xx.config file for the agent. If it refers to tmaitm6/aix513/lib, add the following line in the xx.ipv6config file:

```
export LIBPATH=$CANDLEHOME/tmaitm6/aix523/lib:$LIBPATH
```

If it refers to tmaitm6/aix516/lib, add the following line in the xx.ipv6config file:

```
export LIBPATH=$CANDLEHOME/tmaitm6/aix526/lib:$LIBPATH
```

On Linux, if the agent being configured is based on Tivoli Monitoring V6.1 but installed in a Tivoli Monitoring V6.3 environment, an additional step is required as described in the previous section for AIX.

Check the LD_LIBRARY_PATH variable in the xx.config file for the agent. If it refers to tmaitm6/ls3246/lib, add the following line in the xx.ipv6config file:

```
export LD_LIBRARY_PATH=$CANDLEHOME/tmaitm6/ls3266/lib
```

If it refers to tmaitm6/ls3243/lib, add the following line in the xx.ipv6config file:

```
export LD_LIBRARY_PATH=$CANDLEHOME/tmaitm6/ls3263/lib
```

If it refers to tmaitm6/li6243/lib, add the following line in the xx.ipv6config file:

```
export LD_LIBRARY_PATH=$CANDLEHOME/tmaitm6/li6263/lib
```

On Solaris only, the agent can be configured either for IPv4 or IPv6, but not for both. For IPv4, no action is required. For IPv6, the following line must be added:

```
export KDEB_INTERFACELIST='-'
```

This definition must also be provided in the xx.ipv6config file.

On HP-UX on HP 9000 (PA-RISC) only, the agent can be configured either for IPv4 or IPv6, but not for both. For IPv4, no action is required. For IPv6, the following line must be added:

```
export KDEB_INTERFACELIST='-'
```

This definition must also be provided in the xx.ipv6config file.

Note:

Before starting an agent, be sure that the IPv6 address correctly resolves to the host name because numeric IPv6 addresses are not permitted in managed system names. This rule applies to all distributed agents and to other application agents where the agents provide the managed system name, and the name includes the host name.

The resolved name must be a short name, and not a long fully qualified name, which can cause the managed system name to be truncated, resulting in unpredictable behavior. This rule is not an IPv6-specific constraint.

Warehouse Proxy Agent

You must update the glb_khd.txt file to include the IPv6 protocol when you are configuring the Warehouse Proxy Agent.

The Warehouse proxy agent is configured as described in “Configuring the monitoring agents” on page 444. In addition, the `glb_khd.txt` file must be updated to include the IPv6 protocol. The `glb_khd.txt` file can be found in the `hd/bin` directory under the platform-specific directory in the installation home directory on UNIX, for example: `$CANDLEHOME/li6263/hd/bin` on Linux.

Tivoli Enterprise Portal Server

You configure the Tivoli Enterprise Portal Server for IPv6 communication in the same way that you configure the monitoring agents.

The instructions for “Configuring the monitoring agents” on page 444 also apply to the Tivoli Enterprise Portal Server (or portal server). The product code for the portal server is `cq`.

Adding Application Support to a monitoring server

About this task

If the monitoring server to which application support is to be added is on an IPv6-only system, or configured to use IPv6 only, then the Manage Tivoli Enterprise Management Services console cannot be used to add the application support, and a command-line interface must be used. Use the following procedure to add or remove application support to a management server:

Procedure

1. Locate the executable `kdstsns` in the installed environment. It is usually found under `$CANDLEHOME/$BINARCH/ui/bin` or `$CANDLEHOME/$BINARCH/ms/bin`, where `BINARCH` is the platform architecture designated by IBM Tivoli Monitoring, for example, `li6263` on 32 bit Linux, `aix523` on 32 bit AIX, `hp116` on 64 bit HP-UX, and so on. A complete list of the platforms (`$BINARCH`) and codes can be found in `<ITMHOME>/registry/archdsc.tbl`.
2. Locate the `SQLLIB` directory where the SQL input files for application support reside. This directory is usually under `$CANDLEHOME/tables/cicatrsql`.
3. Change to the `bin` directory under `$CANDLEHOME` directory, for example:
`cd /opt/IBM/ITM/bin`

Create a file called `glb_ipv6tems.txt`, and add a line to it to denote the host on which the monitoring server to be seeded resides, for example:

```
ip6.pipe:atz1004.tivlab.austin.ibm.com
```

4. Create a shell command file called `seedipv6tems.sh` with the following lines (using as an example an installation on 32 bit Linux under `/opt/IBM/ITM`):

```
#!/bin/ksh
export CANDLEHOME=/opt/IBM/ITM
export BINARCH=li6263
export SQLLIB=$CANDLEHOME/tables/cicatrsql/SQLLIB
export KDC_GLBSITES=glb_ipv6tems.txt
$CANDLEHOME/$BINARCH/ui/bin/kdstsns $1 $2
```

Save the file, and make it executable: `chmod +x seedtems.sh`

Note: You can run the following command to find out the **BINARCH** parameter name:

```
$CANDLEHOME/bin/cinfo -s ms | grep BINARCH
```

and look for a line that starts with **ms**, then the **TEMS** name, and then the **BINARCH** parameter name. For example:

```
$CANDLEHOME/bin/cinfo -s ms | grep BINARCH
ms      <TEMS-name>  BINARCH = 1x8266
```

5. Use the following command from the command-line to add support for an application:
`seedipv6tems.sh <app-filename> <TEMS-name>`

where:

app-filename

is the name of the sql file to be added for the desired application.

TEMS-name

is the name of the monitoring server to which support is being added. For example:

```
seedipv6tems kt4.sql TEMSFRT:CMS
```

If multiple applications are to be added, then another script file can be created to run all the commands at once. For example, the file seedipv6all.sh.bat with the following can be used:

```
seedipv6tems.sh kt4.sql TEMSFRT:CMS  
seedipv6tems.sh knt.sql TEMSFRT:CMS
```

6. Use the same procedure to remove application support by using the kxx_del.sql file for application kxx.

Command-line interface (CLI)

The CLI tivcmd and tacmd utilities are used without change in all scenarios.

On AIX only, the following line must be added to the tacmd script file in the \$CANDLEHOME/bin directory:

```
export KDC_FAMILIES='+ip6 +ip4'
```

Configuring Windows monitoring components

The instructions in this section apply to configuring monitoring components on a Windows platform.

Configuring the hub monitoring server

You must first configure hub and remote monitoring servers using the platform-specific installation and configuration tools, which includes selecting a protocol and providing information about the standby hub.

If IPv6 is the preferred protocol for the hub, then the glb_site.txt file must be edited, providing the desired IPv6 family as the protocol. For example:

```
ip6.pipe:proton
```

The glb_site.txt file is found in the cms directory under the installation home directory, which is typically \IBM\ITM.

If the hub has been configured to use a standby server, then the MHM:MIRROR_SITE variable must be edited if the preferred protocol for communication with the standby hub is IPv6. For example:

```
MHM:MIRROR_SITE=ip6.pipe:atlas
```

To set this variable, access the Manage Tivoli Services console, right-click the monitoring server, select **Advanced** → **Edit Variables**, click **Add**, and then add the MHM:MIRROR_SITE variable with the desired value.

If a port other than the default is used, then the KDC_FAMILIES variable must be updated. To set this variable, access the Manage Tivoli Services console, right-click the monitoring server, select **Advanced** → **Edit Variables**, click **Add**, and then add the KDC_FAMILIES variable with the desired value.

```
KDC_FAMILIES=ip6.pipe port:11111 ip6 port:11111 ip6.spipe port:7663 ....
```

If the server has been configured to provide access to other SOAP server hubs, then the kshxhubs.xml file must be updated to enable IPv6 communication with those hubs. This file can be found in the cms\HTML directory under the installation directory, which is typically \IBM\ITM. For example:

```
<?xml version="1.0"?>
<?xml-stylesheet type="text/xsl" href="hubdef.xsl"?>
<ENTERPRISE>
<HUB>
<CMS_Name>
ip6.pipe:saturn[1918]
</CMS_Name>
<Service_Name>newhub</Service_Name>
<Alias>newhub</Alias>
</HUB>
</ENTERPRISE>
```

Configuring the Tivoli Enterprise Monitoring Automation Server

You must change the KDC_FAMILIES variable for the Tivoli Enterprise Monitoring Automation Server to match the updated KDC_FAMILIES variable for the hub monitoring server. In addition, you must also add the automation server connection protocol and port values to the KDC_FAMILIES variable for the automation server.

When updating the KDC_FAMILIES variable for the hub monitoring server for IPv6 communication, as described in “Configuring the hub monitoring server” on page 447, you must apply these same updates to the KDC_FAMILIES variable for the Tivoli Enterprise Monitoring Automation Server.

For example, if the automation server connection protocol and port values are HTTP and 10001, and the KDC_FAMILIES variable for the hub monitoring server was updated to the following:

```
KDC_FAMILIES=ip6.pipe port:11111 ip6 port:11111 ip6.spipe port:7663
```

Then access the Manage Tivoli Services console, right-click the Tivoli Enterprise Monitoring Automation Server, select **Advanced** → **Edit Variables**, click **Add**, and add the KDC_FAMILIES variable with the following value:

```
KDC_FAMILIES=ip6.pipe port:11111 ip6 port:11111 ip6.spipe port:7663 HTTP:10001
```

Note: The automation server will not start if there is an incorrect value for the protocol or port defined in either the KDC_FAMILIES or KDE_TRANSPORT environment variables. Verify that the settings support either HTTP or HTTPS protocol and a port value between 1 and 65535.

Configuring the remote monitoring server

The remote monitoring server uses the entries in the glb_site.txt file to look up the hub. If there is a standby hub, then there are at least two entries in this file.

This file must be edited if the protocol for communicating with one or both the hubs must be set to IPv6. This file is found in the cms directory under the installation home directory, which is typically \IBM\ITM. For example:

```
ip6.pipe:proton
ip.pipe:amoeba
```

If a port other than the default is used, then the KDC_FAMILIES variable must be updated. To set this variable, access the Manage Tivoli Services console, right-click the management server, select **Advanced** → **Edit Variables**, click **Add**, and then add the KDC_FAMILIES variable with the desired value:

```
KDC_FAMILIES=ip6.pipe port:11111 ip6 port:11111 ip6.spipe port:7663 ....
```

Configuring the monitoring agents

Agents use the CT_CMSLIST variable to locate the monitoring server they are required to connect to.

This variable may contain one or more entries, separated by a semi-colon. For example:

```
CT_CMSLIST='ip6.pipe:proton;ip6.pipe:atlas'
```

If IPv6 is the preferred protocol for the agent to communicate with the designated monitoring server, then this variable must be updated. Before doing this, ensure that the monitoring server is configured to communicate using IPv6.

To set this variable, access the Manage Tivoli Services console, right-click the agent, select **Advanced** → **Edit Variables**, click **Add**, and then add the CT_CMSLIST variable with the desired value.

If a port other than the default is used, then the KDC_FAMILIES variable must be updated. To set this variable, access the Manage Tivoli Services console, right-click the agent, select **Advanced** → **Edit Variables**, click **Add**, and then add the KDC_FAMILIES variable with the desired value:

```
KDC_FAMILIES=ip6.pipe port:11111 ip6 port:11111 ip6.spipe port:7663 ....
```

Note:

Before starting an agent, you must ensure that the IPv6 address resolves to the correct host name because numeric IPv6 addresses are not permitted in managed system names. This rule applies to all distributed OS agents and to other application agents where the agents provide the managed system name, and the name includes the host name.

The resolved name must be a short name, and not a long fully qualified name, which can cause the managed system name to truncate, resulting in unpredictable behavior. This rule is not an IPv6-specific constraint.

Warehouse Proxy Agent

You must update the glb_khd.txt file to include the IPv6 protocol when you are configuring the Warehouse Proxy Agent.

The Warehouse proxy agent is configured as described in “Configuring the monitoring agents” on page 448. In addition, the file glb_khd.txt must be updated to include the IPv6 protocol. This file can be found in the TMAITM6 directory under the home directory on Windows (typically \IBM\ITM).

Tivoli Enterprise Portal Server

You configure the Tivoli Enterprise Portal Server for IPv6 communication in the same way that you configure the monitoring agents.

Configuration instructions for agents also apply to the portal server. The product code for the portal server is cq.

Adding Application Support to a monitoring server

If the monitoring server to which application support is to be added is on an IPv6-only system, or configured to use IPv6 only, then the Manage Tivoli Enterprise Management Services console cannot be used to add the application support, and a command-line interface must be used.

About this task

Use the following procedure to add or remove application support to a management server:

Procedure

1. Change to the CNPS directory in the installation home, for example:
`cd \IBM\ITM\cnps`
2. Create a file called glb_ipv6tems.txt, and add a line to it to denote the host on which the monitoring server to be seeded resides, for example:
`ip6.pipe:atz1004.tivlab.austin.ibm.com`

Save the file.

3. Create another file called `seedipv6tems.bat`, with the following lines:

```
set KDC_GLB_SITES=glb_ipv6tems.txt
kdstsns %1 %2
```

If a port other than the default is being used for communication between components, also include the following as the first or the second line, replacing 11111 with the actual port number used in your environment:

```
KDC_FAMILIES=ip6.pipe port:11111 ip6 port:11111
```

Save the file.

4. Use the following command from the command-line to add support for an application:

```
seedipv6tems.sh <app-filename> <TEMS-name>
```

where:

app-filename

is the name of the sql file to be added for the desired application.

TEMS-name

is the name of the monitoring server to which support is being added. For example:

```
seedipv6tems kt4.sql TEMSFRT:CMS
```

If multiple applications are to be added, then another script file can be created to run all the commands at once. For example, the file `seedipv6all.sh.bat` with the following can be used:

```
seedipv6tems.sh kt4.sql TEMSFRT:CMS
seedipv6tems.sh knt.sql TEMSFRT:CMS
```

5. Use the same procedure to remove application support by using the `kxx_del.sql` file for application `kxx`.

Sending files to z/OS systems using FTP

Catalog and attribute files cannot be sent to z/OS systems over IPv6 using the file transfer capability provided by the Manage Tivoli Enterprise Services console provided on the Windows platform.

These files should be transferred using other FTP utilities, such as the one available from the Windows command console.

Tivoli Enterprise Portal browser client

There are no IPv6 configuration requirements for the browser client.

The Internet Explorer browser on Windows 2003 and Windows XP does not allow the use of numeric IPv6 addresses in URLs, so symbolic names must be used to see IPv6 hosts.

Tivoli Enterprise Portal desktop client

Numeric IPv6 can be used to identify the portal server in the desktop client configuration.

If a numeric IPv6 address is used, it must be enclosed in square brackets. For example:

```
[2002:930:9b04:305:9:48:133:98]
```

Command-line interface

The command-line utilities **tivcmd** and **tacmd** can be used without any changes.

Configuring z/OS monitoring components

On the z/OS platform, you can use the ICAT utility for configuring components to use IPv6, as ICAT configuration panels offer IPv6 as one of the protocols.

Manual editing of configuration files is required only if numeric IPv6 addresses must be provided for hosts, or if the IPv4 and IPv6 addresses are bound to two different host names.

The various configuration variables and members that must be updated to provide numeric IPv6 addresses are described in the following sections.

The KDEB_INTERFACELIST_IPv6 parameter in KppENV members in the RKANPARU data set provides a list of interfaces to be used for IPv6 communication. Numeric IPv6 addresses must be enclosed in parentheses. Multiple addresses are separated by spaces. For example:

```
KDEB_INTERFACELIST_IPv6=\
(2002:930:9b04:305:9:48:133:98)
(2002:930:9b04:305:9:48:133:100)
```

The CT_CMSLIST parameter in the KppENV members in the RKANPARU data set provides a list of monitoring servers agents can connect to, in the order provided. For example:

```
CT_CMSLIST=\
IP.PIPE:SP13;\
IP6.PIPE:( 2002:930:9b04:305:9:48:133:98)
```

The KppSSITE member in the RKANPARU data set provides a list of hosts. Numeric IPv6 addresses designating hosts, if provided, must be enclosed in parentheses. For example:

```
IP.PIPE:SP13
IP6.PIPE:( 2002:930:9b04:305:9:48:133:98)
```

The KSHXHUBS member in the RKANPARU data set is an XML document that specifies additional SOAP servers that can be accessed through this hub. If numeric IPv6 addresses are provided, they must be enclosed in parentheses. For example:

```
<?xml version="1.0"?>
<?xml-stylesheet type="text/xsl" href="hubdef.xsl"?>
<ENTERPRISE>
<HUB>
<CMS_Name>
IP6.PIPE:(2002:930:9b04:305:9:48:133:98)[1918]
</CMS_Name>
<Service_Name>newhub</Service_Name>
<Alias>newhub</Alias>
</HUB>
</ENTERPRISE>
```

Firewall considerations

A firewall gateway might be required in some cases.

For information on configuring firewall gateways and operation of other Tivoli Monitoring components in firewall environments, see “Firewall gateway” on page 33.

Chapter 16. Additional Tivoli Enterprise Portal configurations

This chapter discusses how to perform advanced configuration of Tivoli Enterprise Portal components.

- “Connecting the Tivoli Enterprise Portal Server on Windows to a different monitoring server”
- “Using SSL between the portal server and the client” on page 454
- “Configuring historical data collection for the Performance Analyzer warehouse agent” on page 432
- “Reverting from the IBM HTTP Server to the internal web server” on page 456
- “Configuring HTTP communication between the portal client and server” on page 457
- “Firewall network address translation (NAT) or multiple network interface cards” on page 876
- “Firewall scenarios for Tivoli Enterprise Portal” on page 878

It also includes illustrations of firewall scenarios that can help in defining the Tivoli Enterprise Portal Server interface.

Connecting the Tivoli Enterprise Portal Server on Windows to a different monitoring server

When reconfiguring the portal server on Windows for a different Tivoli Enterprise Monitoring Server, a pop-up window is displayed asking if a snapshot of the current Tivoli Enterprise Portal Server data should first be taken.

The default response is **Yes** (see Figure 93).

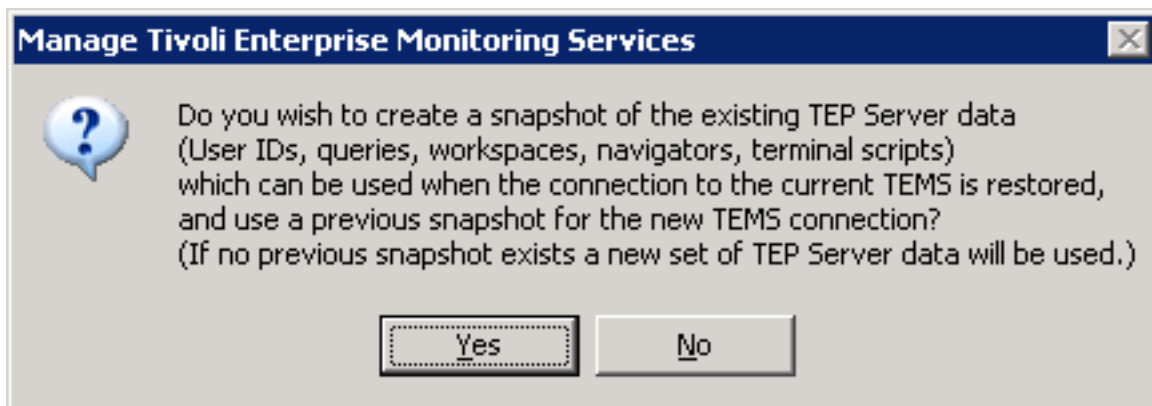


Figure 93. Tivoli Enterprise Portal Server snapshot request screen

Note: When reconfiguring the portal server to communicate with a Hot Standby monitoring server, the correct response is **No**, as the same portal server data should be used for both the primary monitoring server and the secondary monitoring server. See the *IBM Tivoli Monitoring High Availability Guide for Distributed Systems*.

When Yes is selected, a snapshot of the portal server data is taken via the migrate-export process. The data is saved in a file named `saveexport.sql` and is placed in `%CANDLE_HOME%\CNPS\CMS\hostname:port`, where `hostname:port` is the current monitoring server's hostname and connection port number.

When this process completes, the verification screen shown in Figure 94 is displayed.



Figure 94. Tivoli Enterprise Portal Server snapshot verification screen

If no existing snapshot can be found for the monitoring server that is being switched to, a new set of portal server data is created, which means all existing customizations will not be included. To restore these for use with the new monitoring server (if needed), invoke the *migrate-import* process using as input the `saveexport.sql` file created during the previous snapshot request.

When reconfiguring the Tivoli Enterprise Portal Server to switch back to the previous Tivoli Enterprise Monitoring Server, answering **Yes** causes the previous snapshot to be automatically loaded, thus restoring your prior customizations.

Using SSL between the portal server and the client

You can choose to encrypt all communication between the portal server and portal client.

IBM Tivoli Monitoring uses two protocols to provide this level of security between portal server and client server:

- Secure Hypertext Transport Protocol (HTTPS) to retrieve files and Interoperable Object Reference (IOR). The integrated browser in the client provides HTTPS support on the client side.
- Internet Inter-ORB Protocol (IIOP) to secure the communications between the portal server and client. This uses the Secure Sockets Layer (SSL) API provided by VisiBroker. This secure communication uses public key cryptography.

When you install IBM Tivoli Monitoring, the Global Security Toolkit (GSKit) and iKeyman utilities are installed by default on all components. These utilities are used to create and manage the encryption of data between components through the use of digital certificates.

Digital certificates are the vehicle that SSL uses for public-key cryptography. Public-key cryptography uses two different cryptographic keys: a private key and a public key. Public-key cryptography is also known as *asymmetric cryptography*, because you can encrypt information with one key and decrypt it with the complement key from a given public/private key pair.

Public/private key pairs are simply long strings of data that act as keys to a user's encryption scheme. The user keeps the private key in a secure place (for example, encrypted on a computer's hard drive) and provides the public key to anyone with whom the user wants to communicate. The private key is used to digitally sign all secure communications sent from the user; the public key is used by the recipient to verify the sender's signature.

Public/private key pairs are validated by a trusted third party, called a Certificate Authority (CA). An example of a CA is Verisign. If you are setting up your own key pairs, you submit them to the CA for validation.

If you intend to use SSL for communication between the Tivoli Enterprise Portal Server and its clients, use the GSKit provided with IBM Tivoli Monitoring to manage certificates and keys. See the *IBM Tivoli Monitoring Administrator's Guide* for instructions for setting up this encryption.

For additional information about using public/private key pairs, see the iKeyman documentation available at <http://publib.boulder.ibm.com/infocenter/javasdk/v5r0/index.jsp?topic=%2Fcom.ibm.java.security.component.doc.50%2Findex.html>.

Enabling and disabling SSL for the Tivoli Enterprise Portal Server

Secure Sockets Layer (SSL) configuration is for CORBA traffic (as opposed to HTTP traffic) that is used between clients and the portal server. The default port for CORBA traffic is 15001. IBM Tivoli Monitoring is shipped with SSL disabled as the default. If you want to use Secure Sockets Layer communication between the portal server and the portal client, use the steps in this section to enable it.

Note: This procedure disables the primary interface. For instructions on disabling additional interfaces, see Chapter 16, “Additional Tivoli Enterprise Portal configurations,” on page 453.

Windows

1. In the Manage Tivoli Enterprise Monitoring Services window, right-click **Tivoli Enterprise Portal Server**.
2. Click **Advanced** → **Configure TEPS Interfaces**.
3. Highlight **cnps** and click **Edit** in the TEPS Interface Definitions window.
4. Select **Enable SSL for TEP Clients**.
5. Click **OK** to save your changes and close the window.
6. Recycle the service by stopping and starting it.

Linux

1. Change to the *install_dir/bin* directory
2. Run the following command:

```
./itmcmd manage
```
3. In the Manage Tivoli Enterprise Monitoring Services window, right-click **Tivoli Enterprise Portal Server**.
4. Click **Configure**
5. In the first tab, select **Enable SSL for TEP Clients** to enable SSL in the Tivoli Enterprise Portal Server window.
6. Click **OK** to save your changes and close the window.
7. Recycle the service by stopping and starting it.

Disabling SSL

If you do not want to use Secure Sockets Layer communication between IBM Tivoli Monitoring components and the Tivoli Enterprise Portal Server, use the steps outlined in this section to disable it.

Before you begin

Each interface independently enables or disables SSL. If you are using multiple interfaces, you must disable all of them.

Procedure

1. In Manage Tivoli Enterprise Monitoring Services, right-click **Tivoli Enterprise Portal Server**.
2. Click **Advanced** → **Edit ENV** file.
3. Locate the following line:

```
kfw_interface_cnps_ssl=Y
```

4. Change the Y to N.
5. Save the file and exit.
6. Click **Yes** when you are asked if you want to recycle the service.

Reverting from the IBM HTTP Server to the internal web server

The IBM HTTP Server is the default web server used for communication between the portal client and server. This provides increased scalability, performance, and reliability for portal clients. The switch to IBM HTTP Server occurs automatically with no action required on your part.

The IBM HTTP Server does require the following ports to be open in any firewall between the client and the server:

- Port: 15200 – HTTP
- Port: 15201 – HTTPS

If you do not want to open these ports you can use the following sections to revert from the IBM HTTP Server back to the internal web server:

- “Portal server”
- “Browser client”
- “Desktop client”
- “Web Start client” on page 457

Portal server

Take the following steps to configure the Portal client.

Procedure

1. You must add the variable `KFW_IOR_IHS_REDIRECT=N` to the environment file:
 - **Windows** Edit the `install_dir/CNPS/KFWENV` file.
 - **Linux** **AIX** Edit the `install_dir/config/cq.ini` file.
2. Add the variable `KFW_IOR_IHS_REDIRECT=N` to the environment file.

Browser client

Take the following steps to configure the Browser client.

Procedure

1. You must add the command `document.writeln('<PARAM NAME= "cnp.http.useIHS" VALUE="false">');` to the `document.writeln` commands:
 - **Windows** Edit the `install_dir/CNB/applet.html` file.
 - **Linux** **AIX** Edit the `install_dir/platform/cw/applet.html` file.
2. Scroll down to the bottom of the file, where a section of `document.writeln` commands create a set of `<PARAM>` tags. Add the `document.writeln('<PARAM NAME= "cnp.http.useIHS" VALUE="false">');` command.

Desktop client

Take the following steps to configure the Desktop client.

Procedure

1. You must add the JVM parameter `-Dcnp.http.useIHS=false` to the parameter list:
 - **Windows** Edit the `install_dir/CNP/cnp.bat` file and locate the line beginning with `set _CMD=`.

- **Linux** Edit the `install_dir/bin/cnp.sh` file and locate the line beginning with `${TEP_JAVA_HOME}/bin/java`.
2. Add the JVM parameter `-Dcnp.http.useIHS=false` to the parameter list.

Web Start client

Take the following steps to configure the Web Start client.

Procedure

1. Change the JNLP template files:
 - **Windows** Edit the `install_dir/Config/tep.jnlp` file and locate the line beginning with `codebase=`.
 - **Linux** **AIX** Edit the `install_dir/config/tep.jnlp` file and locate the line beginning with `codebase=`.
2. Change the value of the codebase variable to `codebase="http://$HOST$:1920///cnp/kdh/lib/"`. Do not substitute the portal server hostname for the `$HOST$` variable.
3. Locate the **<resources>** section and add the following line directly under the last **<property name...>** line:


```
<property name="cnp.http.useIHS" value="false"/>
```
4. Repeat steps 1 and 2 for the `component.jnlp` file in the specified directory.
5. You must reconfigure the Tivoli Enterprise Portal Server browser client to update all of the `.jnlp` files.

Configuring HTTP communication between the portal client and server

HTTP communication can be used between the portal client and server without the need for the CORBA communication protocol. HTTP communication requires the use of the IBM HTTP Server that is installed automatically with the portal server.

Two new ports are available for client/server communication:

- Port: 15200 – HTTP
- Port: 15201 – HTTPS

Use the following section to change the port numbers used by IHS:

- “Changing the ports used by IHS” on page 459

Use the following sections to switch clients to use HTTP communications with the portal server:

- “Browser client”
- “Desktop client” on page 458
- “Web Start client” on page 458

Browser client

Take the steps outlined in this section to configure the browser client.

Windows

1. In the Manage Tivoli Enterprise Monitoring Services window, right-click **Tivoli Enterprise Portal browser client**.
2. Select **Reconfigure**.
3. In the properties list, locate and double-click the property **tep.connection.protocol**.
4. In the Edit Tivoli Enterprise Portal Browser Parameter window, change the value to **http**.
5. Click **OK**.
6. Click **OK** to close the Configure Tivoli Enterprise Portal Browser window.

1. Edit the `$CANDLEHOME/platform/cw/applet.html` file.
2. Scroll down to the following section at the end of the file:

```
var parameters = { 'cache_archive': cache_archive,
                  'cache_version': cache_version,
                  'java_arguments': javaArgs,
                  'applet_stop_timeout': '5000',
                  'kjr.trace.params': 'ERROR',
                  'cnp.publishurl.encryptuid': 'false',
                  'cnp.cmwapplet.systemout': 'YES',
                  'code': 'candle.fw.pres.CMWApplet',
                  'codebase': 'classes/',
                  'mayscript': 'true',
                  'cache_option': 'plugin',
                  'scriptable': 'true',
                  'kjr.trace.mode': 'LOCAL',
                  'cnp.publish.url': 'true',
                  'cnp.window.timeout': '0',
                  'javax.net.ssl.trustStore': 'jcf.jks'
                };
```

Add the following line: `'tep.connection.protocol': 'http'`

3. To enable secure communication using https, substitute https where http is specified for the new parameter.

For the browser client, the valid options are **iiop** and **http**. These values instruct the client to use either the CORBA IIOP protocol or the HTTP protocol. The use of the secure HTTPS protocol is determined by the protocol specified in the URL used to connect to the portal server. For example, if you connect to the portal server using a URL that starts with **http://**, all communication with the server is unencrypted. The following table describes the browser client behavior:

Table 72. Browser client behavior

tep.connection.protocol value	URL begins with	Result
http	"http://"	All communication uses the HTTP protocol.
https	"https://"	All communication uses the HTTPS protocol.

Desktop client

Take the following steps to configure the desktop client.

Procedure

1. In the Manage Tivoli Enterprise Monitoring Services window, right-click **Tivoli Enterprise Portal desktop client**.
2. Select **Reconfigure**.
3. In the properties list, locate and double-click the property **tep.connection.protocol**.
4. In the Edit Tivoli Enterprise Portal Browser Param window, change the value to **http** or **https**.
5. Click **OK**.
6. Click **OK** to close the Configure Tivoli Enterprise Portal Browser window.

Web Start client

Take the following steps to configure the Web Start client.

Procedure

1. On the host of the Tivoli Enterprise Portal Server, change to the following directory:
 - **Windows** `itinstall_dir\config`. For example, `C:\ibm\itm\config`.
 - **Linux** **UNIX** `itinstall_dir/config`. For example, `/opt/IBM/ITM/config`.
2. Open the `tep.jnlpt` file in a text editor and make the following changes:
 - a. Under the `<resources>` section, add:

```
<property name="tep.connection.protocol" value="http"/>
```
 - b. To use HTTPS protocol, replace the **value** with the string **https**.
3. On the host of the Tivoli Enterprise Portal Server, change to the following directory:
 - **Windows** `%CANDLE_HOME%\CNB\`.
 - **Linux** **UNIX** `$CANDLEHOME/platform/cw/`.
4. Open the `tep.jnlp` file in a text editor and make the following changes:
 - a. Under the `<resources>` section, add the following parameter to the list of JVM parameters:

```
<property name="tep.connection.protocol" value="http"/>
```
 - b. To use HTTPS protocol, replace the **value** with the string **https**.
5. You must reconfigure the browser client for these changes to take effect.

Changing the ports used by IHS

Complete the steps outlined in this section to change the ports used by IHS to listen for HTTP/S communication between the client and server.

1. On the host of the Tivoli Enterprise Portal Server, change to the following directory:
 - **Windows** `itinstall_dir\IHS\CONF`
 - **Linux** **UNIX** `itinstall_dir/arch/iu/ihs/HTTPServer/conf`
2. Open the `httpd.conf` file in a text editor and change the following listen directive:

```
Listen 15200
```

```
ServerName <hostname>:15200
```

to:

```
Listen <http port #>
```

```
ServerName <hostname>:<http port #>
```

where **<http port #>** is the new port to use for HTTP communication.

3. To change the HTTPS port, change the following listen directive:

```
Listen 0.0.0.0:15201
```

to:

```
Listen 0.0.0.0:<https port #>
```

where **<https port #>** is the new port to use for HTTPS communication.

4. Change to the following directory:
 - **Windows** `itinstall_dir\CNPSJ\profiles\ITMProfile\config\cells\ITMCell`
 - **Linux** **UNIX** `itinstall_dir/arch/iw/profiles/ITMProfile/config/cells/ITMCell`
5. Open the `virtualhosts.xml` file in a text editor and change the following alias attribute:

```
<aliases xmi:id="HostAlias_1354390002404" hostname="*" port="15200"/>
```

to:

```
<aliases xmi:id="HostAlias_1354390002404" hostname="*" port="<http port #>"/>
```
6. To change the HTTPS port, change the following alias definition:

```
<aliases xmi:id="HostAlias_1354390002524" hostname="*" port="15201"/>
```

to:

```
<aliases xmi:id="HostAlias_1354390002524" hostname="*" port="<https port #>"/>
```

7. Restart the portal server to implement these changes.

For the Desktop client and Web Start client, you must explicitly define the new port number. The browser client determines the port based on the URL used to launch the client.

For the desktop client, take the following steps to change the port:

1. Open the client launcher in a text editor:
 - **Windows** Edit `itminstall_dir\CNP\cnp.bat`
 - **Linux** **UNIX** Edit `itminstall_dir/arch/cj/bin/cnp.sh`
2. Locate the line beginning with:
 - **Windows** `"set _CMD="`
 - **Linux** **UNIX** `${TEP_JAVA_HOME}`
3. Add the following parameter to the command-line. If the Tivoli Enterprise Portal is configured to use the HTTP or HTTPS protocol to communicate with the Tivoli Enterprise Portal Server, use the HTTP/S parameter. Otherwise, use the IIOP parameter:
 - **IIOP:** `-Dcnp.http.url.port=<port #>`
 - **HTTP/S:** `-Dtep.connection.protocol.url.port=<port #>`

where **<port #>** is the new port to use for communication with the server.

For the Web Start client, take the following steps to change the port:

1. On the host of the Tivoli Enterprise Portal Server, change to the following directory:
 - **Windows** `itminstall_dir\config`. For example, `C:\ibm\itm\config`
 - **Linux** **UNIX** `itminstall_dir/config`. For example, `/opt/IBM/ITM/config`
2. Open the `tep.jnlpt` file in a text editor and make the following changes. If the Tivoli Enterprise Portal is configured to use the HTTP or HTTPS protocol to communicate with the Tivoli Enterprise Portal Server, use the HTTP/S parameter. Otherwise, use the IIOP parameter:
Under the `<resources>` section, add:
 - **IIOP:** `<property name="cnp.http.url.port" value="<port #>">`
 - **HTTP/S:** `<property name="tep.connection.protocol.url.port" value="<port #>"/>`

where **<port #>** is the new port to use for communication with the server.

3. You must reconfigure the Tivoli Enterprise Portal Server for these changes to take effect.

Chapter 17. Configuring IBM Tivoli Monitoring Web Services (the SOAP Server)

IBM Tivoli Monitoring Web Services provide an industry-standard open interface into products that use the Tivoli Management Services framework. This open interface provides easy access to performance and availability data, allowing you to use this information for advanced automation and integration capabilities.

Web Services implements a client/server architecture. The client sends SOAP requests to the SOAP server. The server receives and processes the SOAP requests from the client. Predefined SOAP methods let you perform many functions within the monitored environment. Using Web Services requires a basic understanding of SOAP, XML and XML Namespaces, and the Web Services Description Language (WSDL).

By default, the SOAP server is enabled on all hub monitoring servers. This chapter describes how to configure security on a SOAP server, and how to configure communication between SOAP servers.

Note: You cannot make SOAP requests from IBM Tivoli Monitoring to earlier SOAP servers (such as those in OMEGAMON Platform V350/360).

For complete information about using Web Services and customizing the SOAP interface for your site, see the *IBM Tivoli Monitoring Administrator's Guide*.

Table 73 outlines the steps required to configure the SOAP server.

Table 73. Overview of SOAP Server configuration steps

Task	Where to find information
Define the hubs with which your SOAP Server communicates.	"Defining hubs"
Create users and grant them access.	"Adding users" on page 464
Verify that you have successfully configured SOAP.	"Verifying the configuration" on page 466

Defining hubs

In this step you activate a SOAP server and define hubs with which it communicates.

You can use Manage Tivoli Enterprise Monitoring Services to configure the SOAP server. On Linux and UNIX you can also use the **itmcmd config** command. After you configure the SOAP Server, you must restart the Tivoli Enterprise Monitoring Server.

- "Windows: Defining hubs"
- "UNIX and Linux: Defining hubs (GUI procedure)" on page 462
- "UNIX and Linux: Defining hubs (CLI procedure)" on page 463

Windows: Defining hubs

Define the SOAP server hubs on a Windows system.

About this task

Use the following steps to define SOAP hubs on Windows:

Procedure

1. Start Manage Tivoli Enterprise Monitoring Services (select **Start → (All) Programs → IBM Tivoli Monitoring → Manage Tivoli Monitoring Services**).
2. In the Manage Tivoli Enterprise Monitoring Services window, right-click **Tivoli Enterprise Monitoring Server**.
3. Click **Advanced → Configure SOAP Server Hubs**. The SOAP Server Hubs Configuration window is displayed. If the name of the local hub does not appear in the tree, define the local hub, including assigning user access, before defining the hubs with which it communicates.
4. Click **Add Hub**. The Hub Specification window is displayed.
5. Select the communications protocol to be used with the hub from the **Protocol** menu.
6. Specify an alias name in the **Alias** field. The alias for the local hub monitoring server must always be "SOAP". For hubs with which the local SOAP Server communicates, you may choose any alias (for example, HUB2). Alias names can be a minimum of 3 characters and a maximum of 8 characters.
7. Do one of the following:
 - If you are using TCP/IP or TCP/IP Pipe communications, complete the fields in Table 74:

Table 74. TCP/IP Fields in Hub Specification dialog

Field	Description
Hostname or IP Address	The host name or TCP/IP address of the host computer.
Port	The TCP/IP listening port for the host computer.

- If you are using SNA communications, complete the fields in Table 75:

Table 75. SNA Fields in Hub Specification dialog

Field	Description
Network Name	Your site SNA network identifier.
LU Name	The LU name for the monitoring server. This LU name corresponds to the Local LU Alias in your SNA communications software.
LU6.2 LOGMODE	The name of the LU6.2 logmode. Default: CANCTDCS.
TP Name	The Transaction Program name for the monitoring server.

8. Click **OK**.

Results

The server tree is displayed, with the newly defined hub. You can define any hubs with which the local SOAP Server will communicate by repeating steps 4 - 7, or you can specify which user IDs can access the SOAP Server you just defined and what access privileges they have. For more information, see "Adding users" on page 464.

UNIX and Linux: Defining hubs (GUI procedure)

Define the SOAP server hubs on a Linux/UNIX system using the GUI procedure.

About this task

Use the following steps to define SOAP hubs on UNIX or Linux using Manage Tivoli Enterprise Monitoring Services:

Procedure

1. Change to the *itm_install_dir/bin* directory and start Manage Tivoli Enterprise Monitoring Services by entering the following command:

```
./itmcmd manage
```

The Manage Tivoli Enterprise Monitoring Services window is displayed.

2. Right-click **Tivoli Enterprise Monitoring Server** and select **Configure** from the popup menu. The Configure TEMS window is displayed.
3. Click **Save**. The SOAP Server Hubs Configuration window "UNIX and Linux: Defining hubs (GUI procedure)" on page 462 is displayed.
4. Confirm that the host name or IP address, port number, and protocol displayed for the hub monitoring server are correct. If not, correct them. If the name of the local hub does not appear in the tree, define the local hub before defining the hubs with which it communicates. The alias for the local hub must always be "SOAP".
5. To add another hub:
 - a. Type the name or IP address and the port number of the host in the appropriate fields. The port should be match the hub's protocol port number. By default, this is 1918.
 - b. Specify an alias name in the **Alias** field. Alias names can be a minimum of 3 characters and a maximum of 8 characters (for example, HUB2).
 - c. Select the communications protocol to be used with the hub from the **Transport** menu.
6. Click **Add Host**. The server tree is displayed, with the newly defined hub.

Results

You can now define any hubs with which the local SOAP Server will communicate by repeating step 5, or you can specify which user IDs can communicate with the local SOAP Server and what access privileges they have.

UNIX and Linux: Defining hubs (CLI procedure)

Define the SOAP server hubs on a Linux/UNIX system using the CLI procedure.

About this task

Complete the following steps to configure the SOAP server:

Procedure

1. On the host of the hub monitoring server on which you want to configure Web Services, change to the *install_dir/bin* directory and enter the following command:
2. Accept the default values, which should reflect the choices made during the last configuration, until you see the following prompt:

```
*****
```

```
Editor for SOAP hubs list
```

```
*****
```

```
Hubs
```

```
## CMS_Name
```

```
1 ip.pipe:TEMS_NAME[port_#]
```

```
1)Add, 2)Remove ##, 3)Modify Hub ##, 4)UserAccess ##, 4)Cancel, 5)Save/exit:
```

3. To add a hub with which the local hub can communicate:

- a. Type 1 and press Enter.
- b. Respond to the following prompts as shown in :
 - Network Protocol [ip, sna, ip.pipe, or ip.spipe] (Default is: ip):
 - CMS Name (Default is: *local_host*):
 - Port Number (Default is: 1918):
 - Alias (Default is: SOAP):

Table 76. SOAP hub configuration values

Configuration	Value
Network Protocol	The communications protocol configured for the hub monitoring server
CMS Name	The host name of the hub monitoring server. The host name must be fully
Port Number	The protocol port for the hub monitoring server.
Alias	An alias for the hub. Alias names can be a minimum of 3 characters and a maximum of 8 characters. The alias for the local hub must always be SOAP, but you must create a new alias for additional hubs (for example, HUB2).

After you enter the alias, the list of hubs is displayed with the new hub added. For example,

Hubs

```
##      CMS_Name
1      ip.pipe:chihuahua[1918]
2      ip:maple[1918]
```

1)Add, 2)Remove ##, 3)Modify Hub ##, 4)UserAccess ##, 5)Cancel, 6)Save/exit:

You can continue to add hubs, or you can proceed to define user access for the hubs you have already defined.

Adding users

Access to SOAP server can be secured in one of two ways: by enabling **Security: Validate User** and creating user accounts on the host of the hub monitoring server, or by adding specific users to the server definition.

If **Security: Validate User** is *not* enabled and no users are added to the server definition, the SOAP server honors all requests regardless of the sender. If **Security: Validate User** is enabled on the hub monitoring server, the SOAP server honors requests only from users defined to the operating system or security authorization facility of the monitoring server host.

Note: When you enable the **Security: Validate Users**, all SOAP requests are authenticated except for CT_EMail and CT_Export requests. If you want the SOAP server to validate users of those requests, you must also set the SOAP_IS_SECURE environment variable for the hub monitoring server.

However, if any users are added to the SOAP server definition, only those users who have also been defined to the operating system or the security authorization facility of the monitoring server host have access to the server, regardless of whether or not **Security: Validate User** is enabled.

In this step you define users to a SOAP Server and specify the access privileges for each user: Query or Update. The access privileges control what methods the user is allowed to use. Update access includes Query access. Table 77 lists the methods associated with each access. For information on using these methods, see the *IBM Tivoli Monitoring Administrator's Guide*.

Table 77. SOAP methods associated with access privileges

Method	Update	Query
CT_Acknowledge	x	

Table 77. SOAP methods associated with access privileges (continued)

Method	Update	Query
CT_Activate	x	
CT_Alert	x	
CT_Deactivate	x	
CT_Email	x	x
CT_Execute	x	x
CT_Export	x	x
CT_Get	x	x
CT_Redirect	x	x
CT_Reset	x	
CT_Resurface	x	
CT_Threshold	x	
WTO	x	x

After you configure the SOAP Server, you must restart the Tivoli Enterprise Monitoring Server.

Windows: Adding users

Add a user to a SOAP server on Windows.

About this task

Use the following steps to define users to a SOAP server:

Procedure

1. In the SOAP Server Hubs Configuration window, select the server (click anywhere within the server tree displayed).
2. Under Add User Data, type the user name. If **Security: Validate User** is enabled, user IDs must be identical to those specified for monitoring server logon validation. Access is restricted to only that monitoring server to which a user has access.
3. Click the type of user access: **Query** or **Update**.
4. Click **Add User**. The server tree is updated, showing the user and type of access.
5. To delete a user: Select the user name from the tree and click **Delete Item**.
6. To delete a hub: Click anywhere within the hub's tree and click **Clear Tree**.

UNIX or Linux: Adding users (GUI)

Add a user to a SOAP server on Linux/UNIX systems using the GUI method.

About this task

Complete the following steps to define which users can make requests to a hub and assign access privileges to those users, using a graphical user interface:

Procedure

1. If you are not already in the SOAP Server Hubs Configuration window, perform steps 1 - 3 in "UNIX and Linux: Defining hubs (GUI procedure)" on page 462.
2. In the hub tree, select the **Access** node for the hub to which you want to define users. The values for that hub are displayed in the section above the tree.

3. Type the user ID you want to add in the **User** field. If the **Security: Validate User** option is enabled on the hub, the user ID must be a valid user ID on the hub system.
4. Use the radio buttons in the **Access** field to select the type of access you want to grant to the user. See Table 77 on page 464 for a list of the method each type of access includes.
5. Click **Add User**. The user ID appears under the appropriate subnode of the Access node of the selected hub.
6. To add additional users, repeat steps 2 through 5.
7. Click **OK** to close the window and save the changes.

UNIX or Linux: Adding users (CLI)

Add a user to a SOAP server on Linux/UNIX systems using the CLI method.

About this task

Complete the following steps to define users to a hub and assign access privileges using the command-line interface:

Procedure

1. If you are not already at the following prompt, perform steps 1 and 2 in “UNIX and Linux: Defining hubs (CLI procedure)” on page 463:

```
*****
```

```
Editor for SOAP hubs list
```

```
*****
```

```
Hubs
```

```
## CMS_Name
```

```
1 ip.pipe:TEMS_NAME[port_#]
```

```
1)Add, 2)Remove ##, 3)Modify Hub ##, 4)UserAccess ##, 4)Cancel, 5)Save/exit:
```

2. To define users and assign access privileges, enter 4 followed by a space, and then the number (in the list) of the hub you want to configure. For example:

```
1)Add, 2)Remove ##, 3)Modify Hub ##, 4)UserAccess ##, 5)Cancel, 6)Save/exit:4 1
```

The following prompt is displayed:

```
=> Query Access:
```

```
=> Update Access:
```

```
1)QueryAdd <user>, 2)UpdateAdd <user>, 3)Remove ##, 4)Exit :
```

Any users already defined are listed under the corresponding access.

3. To define a user with Query access, enter 1 followed by a space and the user ID. To define a user with Update access, enter 2 followed by a space and the user ID. See Table 77 on page 464 for a list of the methods associated with each type of access.

Note: If the **Security: Validate User** option is enabled, the user ID must be a valid user on the hub system.

The prompt appears again with the new user added to the appropriate access list. You can continue to add users by repeating step 3 or complete the configuration by typing 4 and pressing Enter.




Verifying the configuration

Verify that you have configured the SOAP server correctly.

About this task

In this step you verify that SOAP has been correctly configured by starting the SOAP client and making a request using the command-line utility kshsoap.

Procedure

1. Verify that the monitoring server that you used to enable SOAP is running. If not, start it.
2. Open a command window.
3.  Change to the *install_dir*\cms directory or   The *install_dir*/TBD.
4. In the current directory, create an ASCII text file named SOAPREQ.txt that contains the following SOAP request:

```
<CT_Get><object>ManagedSystem</object></CT_Get>
```

Or if security has been enabled:

```
<CT_Get><userid>logonid</userid><password>password</password> \
<object>ManagedSystem</object></CT_Get>
```

5. Create another ASCII text file named URLS.txt that contains URLs that will receive the SOAP request. For example: `http://hostname:1920//cms/soap`

6. Run the following command:

```
kshsoap SOAPREQ.txt URLS.txt
```

SOAPREQ.txt is the name of the file that contains the SOAP request and URLS.txt is the name of the file that contains the URLs.

Results

The kshsoap utility processes the SOAPREQ.txt file and displays the output of the SOAP request in the command window. The SOAP request is sent to each URL listed in URLS.txt, and the SOAP response from each URL displays in the DOS window.

Chapter 18. Performance tuning

This chapter contains information about optimizing the performance of several components within an IBM Tivoli Monitoring environment.

The following topics are considered:

- “Disabling TCP-delayed acknowledgments on AIX systems”
- “Tivoli Enterprise Monitoring Server” on page 470
- “Tivoli Enterprise Monitoring agents” on page 472
- “Agentless Monitoring” on page 473
- “Tivoli Enterprise Portal Server” on page 474
- “Tivoli Enterprise Portal client” on page 478
- “Tivoli Data Warehouse” on page 480
 - “Historical data collection” on page 481
 - “Warehouse Proxy Agent” on page 482
 - “Summarization and Pruning Agent” on page 486
 - “Database tuning” on page 489
- “Optimizing queries” on page 502
- “Optimizing situations” on page 504

This chapter includes some sections from the redbook *IBM Tivoli Monitoring: Implementation and Performance Optimization for Large Scale Environments*.

Appendix E, “Agent configuration and environment variables,” on page 887 provides an extensive list of environment variables that can be customized for different components in the Tivoli Monitoring environment. The following sections highlight specific environment variables to consider in performance tuning of the Tivoli Monitoring environment.

Note: When changing environment variables in configuration files, note that whenever maintenance or reconfiguration takes place in your environment, these changes might be lost and need to be reapplied.

Disabling TCP-delayed acknowledgments on AIX systems

On AIX systems, the default behavior for TCP connections results in delayed acknowledgments (Ack packets). When `tcp_nodelayack` is set to 0 (the default setting), TCP delays sending Ack packets by up to 200ms, the Ack attaches to a response, and system overhead is minimized.

Setting the `tcp_nodelayack` parameter to 1 causes TCP to send immediate acknowledgment (Ack) packets to the sender.

Setting `tcp_nodelayack` to 1 will cause slightly more system overhead, but can result in much higher performance for network transfers if the sender is waiting on the acknowledgment from the receiver.

Measurements of communication between Tivoli Monitoring components have shown that setting `tcp_nodelayack` to 1 can significantly improve performance.

To make the parameter setting, issue the following:

```
# no -p -o tcp_nodelayack=1
Setting tcp_nodelayack to 1
Setting tcp_nodelayack to 1 in nextboot file
```

The -p flag makes the change persistent, so that it is still in effect at the next boot. This is a dynamic change that takes effect immediately.

Tivoli Enterprise Monitoring Server

This section provides information about parameters you might consider editing to improve either hub or remote monitoring server performance.

The parameters are set in the following files according to operating system:

Windows

ITM_HOME/cms/KBBENV

For example: C:\IBM\ITM\cms\KBBENV

Linux

UNIX

ITM_HOME/config/*tems_hostname_ms_tems_name*.config

For example: /opt/IBM/ITM/config/edinburg_ms_labtems.config

z/OS

&shilev.&rtename.RKANPARU(KDSENV)

For example: ITM.SYP1.RKANPARU(KDSENV)

Note: The &shilev and &rtename are variables that correspond to high-level qualifiers of the RKANPARU(KDSENV) partitioned data set. These variables can take 1 - 8 characters.

On each occasion maintenance or reconfiguration takes place in your environment these files might be recreated and changes lost and need to be reapplied.

The following list contains the settings that might affect the monitoring server performance:

CMS_DUPER

This parameter enables or disables situation synchronization of common filter objects monitored by agents or endpoints. Enabling this setting in monitoring server environments with predominantly z/OS address space applications for example, OMEGAMON XE for CICS or Sysplex, improves performance and response time by limiting data collection samplings on behalf of running situations. Enable it by setting the value to YES. Disable by setting the value to NO. This parameter is enabled by default.

EVENT_FLUSH_TIMER

This parameter is set (in minutes) to set an interval at which time pending I/Os are committed to situation status history as a background writer and garbage collection task. This feature improves performance of incoming event throughput into the hub monitoring server per arriving situation statuses.

EIB_FLUSH_TIMER

This parameter is used to specify in minutes how long the monitoring server waits before resetting distribution and database event requests to an initial state. This frees held resources by the request if no event information has been able to get processed in the specified time. The default setting is 2 minutes. If event requests do not respond within 2 minutes you might consider allocating a higher minutes setting to allow requests more time to process. This is particularly valid in larger, more complex, environments.

DELAY_STOP

This parameter is used to specify in seconds how long to delay monitoring server shutdown for UNIX and Linux monitoring servers, as invoked by the itmcmd server stop *tems_name* command. The default value is 60 seconds. The delay is used to allow network connections to close before an immediate restart of the monitoring server with the itmcmd server start *tems_name* command.

If you do not immediately restart the monitoring server after shutting it down, this parameter can be set to a lower value to cause the monitoring server shutdown to complete more quickly.

KGLCB_FSYNC_ENABLED

This parameter is not available on Windows systems.

For Linux and UNIX platforms, this variable can be used to specify whether the `fsync()` system call should be invoked after writes to the file system. You can set this configuration variable in the standard configuration file for the monitoring server. By default, for maximum reliability, `fsync()` is called. If, and only if, the following line is added to the monitoring server configuration file, `fsync()` calls are omitted:

```
KGLCB_FSYNC_ENABLED='0'
```

The default behavior is to call `fsync()` after writes, which is equivalent to the setting:

```
KGLCB_FSYNC_ENABLED='1'
```

The `fsync()` system call flushes the dirty pages from the file system to disk and protects against loss of data in the event of an operating system crash, hardware crash, or power failure. However, it can have a significant negative effect on performance because in many cases it defeats the caching mechanisms of the platform file system. On many UNIX platforms, the operating system itself syncs the entire file system on a regular basis. For example, by default the synced daemon that runs on AIX syncs the file system every 60 seconds. This limits the benefit of `fsync()` calls by application programs to protecting against database corruption in the most recent 60 second window.

Tivoli Enterprise Monitoring Server tuning recommendations for large-scale environments

This section provides tuning recommendations you might consider to improve either hub or remote monitoring server performance in large scale environments.

Review these topics and considerations for relevancy to your environment.

Avoid distributing unused situations

To reduce Tivoli Enterprise Monitoring Server memory usage and minimize the amount of CPU processing when many agents connect, consider reducing the total number of situations distributed by avoiding distribution of situations that are not being used.

Some situations, including predefined situations, have the default distribution set as a managed system group. These situations are distributed to all managed Tivoli Enterprise Monitoring Servers in the managed system group, even if the situation is not being used. Limiting the distribution to only the managed Tivoli Enterprise Monitoring Server where the situation is used minimizes the total number of situations distributed from the remote Tivoli Enterprise Monitoring Server, and minimizes the CPU processing when many agents connect. The distribution specification for a situation can be changed using the Situation Editor or the `tacmd editsit` command.

Check ulimit settings for open file descriptors (UNIX/Linux only)

The Tivoli Enterprise Monitoring Server can use many file descriptors, especially in a large environment. On UNIX and Linux Tivoli Enterprise Monitoring Servers, the maximum number of file descriptors available to a process is controlled by user limit parameters. To display the current user limits, use the `ulimit -a` command.

The `nfiles` parameter is the number of file descriptors available to a process. When IP:PIPE or IP:SPIPE are used for agent connectivity, persistent TCP connections are maintained to each agent, and each connection requires a file descriptor. If a file descriptor is not available when needed, unexpected behavior can occur, including program failures. For the Tivoli Enterprise Monitoring Server process (`kdsmain`), the `nfiles` parameter should be set larger than the maximum number of agent connections that will be maintained. Consider increasing the value to 8000 file descriptors or more.

There are other user limit parameters that control how much data, stack, and memory are available to a process. For large environments, consider increasing these memory-related user limit parameters for the monitoring server (kdsmain) process using the `ulimit` settings.

The method for changing the user limit parameters is operating system specific. Consult the operating system manuals for information about how to configure the user limit parameters.

Disable extra `fsync()` calls by specifying `KGLCB_FSYNC_ENABLED=0` (UNIX/Linux only)

The `KGLCB_FSYNC_ENABLED` parameter specifies whether the `fsync()` system call should be invoked after writes to the file system. The default value is 1, causing `fsync()` to be called after every write for maximum reliability. When `KGLCB_FSYNC_ENABLED=0` is set, `fsync()` calls are omitted after every write.

For more information, see the parameters listed in “Tivoli Enterprise Monitoring Server” on page 470.

Minimize Tivoli Enterprise Monitoring Server shutdown time by specifying `DELAY_STOP=1` (UNIX/Linux only)

The `DELAY_STOP` parameter specifies in seconds how long to delay monitoring server shutdown, as invoked by the `itmcmd server stop <Tivoli Enterprise Monitoring Server_name>` command.

For more information, see the parameters listed in “Tivoli Enterprise Monitoring Server” on page 470.

Avoid storing short-term history data at the Tivoli Enterprise Monitoring Server, if possible

For best performance, when configuring historical data collection, store the short-term history data at the agent rather than the Tivoli Enterprise Monitoring Server when possible. Remember that for some agent types, data is required to be stored at the Tivoli Enterprise Monitoring Server.

Storing short-term history data at the agent reduces usage on the Tivoli Enterprise Monitoring Server to collect and store data for all of the connected agents. When an attribute group has short term-history stored at the Tivoli Enterprise Monitoring Server, data from all the agents is stored in a single short-term history file. When a request is made for short-term history for an agent, the Tivoli Enterprise Monitoring Server must read all of the agent data stored in the history file to satisfy the request. This process is much less efficient. Storing short-term history data at the Tivoli Enterprise Monitoring Server increases the Tivoli Enterprise Monitoring Server processing usage, and lowers the number of agents that can be supported by the Tivoli Enterprise Monitoring Server.

Tivoli Enterprise Monitoring agents

This section describes agent environment variables to consider when tuning your Tivoli Monitoring environment.

CTIRA_RECONNECT_WAIT

Time interval, in seconds, for the agent to wait between attempts to register with a Tivoli Enterprise Monitoring Server. The default value is 600 seconds. Consider setting this value to a value equivalent to the `CTIRA_HEARTBEAT` setting, which is specified in minutes. For example, if the `CTIRA_HEARTBEAT` value has been set at 3 minutes, consider setting `CTIRA_RECONNECT_WAIT` to 180 seconds.

CTIRA_MAX_RECONNECT_TRIES

Number of consecutive times without success the agent attempts to connect to a Tivoli Enterprise Monitoring Server before giving up and exiting. The default value is 720. Using the default value along with the default `CTIRA_RECONNECT_WAIT` setting, the agent attempts to connect to the Tivoli Enterprise Monitoring Server for 432000 seconds (5 days) before giving up and exiting. If you lower the value for `CTIRA_RECONNECT_WAIT`, consider increasing this value to maintain

an equivalent retry period. For example, if you lower the CTIRA_RECONNECT_WAIT value to 180 seconds, consider increasing this value to 2400.

Agentless Monitoring

This section provides tuning recommendations you might consider to improve Agentless Monitoring performance in large-scale environments.

Review these topics and considerations for relevancy to your environment.

Adjust the thread pool size, refresh interval, and cache time-to-live parameters for your environment

The main thread pool parameters are CDP_DP_THREAD_POOL_SIZE and CDP_DP_REFRESH_INTERVAL.

For a complete list of Agentless Monitoring parameters, see the *Agent Builder User's Guide*. The following are some of the main parameters:

CDP_DP_THREAD_POOL_SIZE

Controls the number of threads in the thread pool.

CDP_DP_REFRESH_INTERVAL

Controls how often the thread pool collects the data for each attribute group. The refresh interval should also be correlated with CDP_DP_CACHE_TTL. Both are specified in seconds, and default to 60. If the refresh interval was stretched out to, for example, 300 while the cache ttl stayed at 60, then there would be a 4 minute period during which the cache was considered expired, and any queries that came in during that window would cause an on-demand collection to happen. The best practice is to keep these two values the same.

To determine if the value for CDP_DP_THREAD_POOL_SIZE is correct, check the Thread Pool Status attribute group and look at the Average Queue Length and Average Job Wait fields, and the Intervals Skipped field of the Performance Object status. If these numbers are getting large, and the processor load on the system is in the acceptable range, then the size of the thread pool can be increased. One thread per subnode might be a good starting point, although if there are fewer than 60 nodes, there is probably no reason to decrease below the default 60 threads. If the load on the system is too high, or if the Average Queue Length, Average Job Wait, and Intervals Skipped do not decrease, collecting the data less often can decrease the load. The default refresh interval is 60 seconds, so increasing that number, and the CDP_DP_CACHE_TTL as well, can decrease the overall load.

To monitor many monitoring servers, spread the target monitoring servers across multiple process instances

Although it is possible to monitor 100 target monitoring servers from a single process instance, if historical data collection is used the attribute group data from all of the target monitoring servers is written to a single attribute group file.

This exhibits the same behavior as keeping agent historical data on the monitoring server rather than the monitoring agent. When a Tivoli Enterprise Portal client requests short-term historical data from a single target system, the entire attribute group file must be ready to satisfy the request.

To monitor a large number of target monitoring servers using agentless monitoring, it is better to use multiple process instances on the same system, and spread out the target monitoring servers across the process instances. This keeps the number of target monitoring servers with data in a single file to a smaller number, and improves response time for Tivoli Enterprise Portal requests for target system short-term historical data.

Avoid collecting historical data for the KR2_WMI_Event_Log attribute group

The KR2_WMI_Event_Log attribute group generates a high number of rows and collecting historical data for this attribute group can affect performance.

Tivoli Enterprise Portal Server

The Tivoli Enterprise Portal Server (portal server) acts as a conduit for Tivoli Enterprise Portal clients requesting data for analysis from monitoring agents and other components within the enterprise. The portal server connects directly to the hub monitoring server, which it queries for enterprise information and receiving updates as they occur. As it is responsible for handling large amounts of data, it can be a potential bottleneck within the IBM Tivoli Monitoring environment. This section outlines considerations to optimize the portal server performance.

The portal server database stores information related to the presentation of monitored data at the portal client, including definitions related to users, workspaces and views. After IBM Tivoli Monitoring V6.1 Fix Pack 3, the portal server also stores information related to events and event attachments in the portal server database. Before IBM Tivoli Monitoring V6.1 Fix Pack 3, the portal server database required little or no tuning. In environments with a moderate to high rate of events, the portal server database might require some tuning to optimize performance. In particular, the KFWTSIT table, which is used to store events, can grow large.

If you are using DB2 for the portal server database, consider the following tuning recommendations:

- Manual control of the DB2 bufferpool size for the TEPS database is recommended over letting DB2 self-tune the bufferpool size. If DB2 is configured to do self-tuning of the bufferpool size, the size chosen is often much larger than necessary for good Tivoli Enterprise Portal Server performance. A bufferpool size of 10,000 4K pages (about 40 MB) should be sufficient for most environments. The default buffer pool size is 250 4K pages. Increase the size of this buffer pool to minimize disk I/O activity.
- The following commands illustrate how to increase the size of the buffer pool to 10000 4K pages (from a DB2 command prompt):

—

```
CONNECT TO TEPS;  
ALTER BUFFERPOOL IBMDEFAULTBP IMMEDIATE SIZE 10000;  
CONNECT RESET;
```

- Because the KFWTSIT table is the most active table, use the runstats and reorg facilities on a regular (for example, daily) basis.
- The following commands illustrate how to run the RUNSTATS facility on the KFWTSIT table (from a DB2 command prompt):

—

```
CONNECT TO TEPS;  
RUNSTATS ON TABLE TEPS.KFWTSIT AND INDEXES ALL ALLOW WRITE ACCESS ;  
CONNECT RESET;
```

- The following commands illustrate how to run the REORG facility on the KFWTSIT table (from a DB2 command prompt):

—

```
CONNECT TO TEPS;  
REORG TABLE TEPS.KFWTSIT INPLACE ALLOW WRITE ACCESS START ;  
CONNECT RESET;
```

Note: These tuning changes can also be made using the DB2 Control Center.

Portal server parameter tuning

This section provides some information about parameters you might consider editing to improve portal server performance.

The parameters are set in the following files according to operating system:

Windows

`ITM_HOME/CNPS/kfwenv`

For example: `C:\IBM\ITM\CNPS\kfwenv`

Linux

UNIX

`ITM_HOME/config/cq.ini`

For example: `/opt/IBM/ITM/config`

On each occasion maintenance or reconfiguration takes place in your environment these files might be recreated and changes lost and must be reapplied.

Note: For parameter changes made in the portal server configuration file to take effect, the portal server must be stopped and restarted.

KFW_CMW_EVENT_SLEEP

In a complex environment, you might have a number of events occurring simultaneously. This variable specifies a time in seconds the portal server waits between processing similar events. Consider setting this variable to a value of less than 10 seconds if you are experiencing slow performance, such as portal client refresh, as a result.

KFW_EVENT_ASSIST

The Event Assistant is an internal component within the portal server that allows the user to:

- Associate attachments to events, such as logs
- Assign ownership to events and transfer ownership between users
- View events specific to the current user for the convenience of working with events which they own
- View closed events along with any associated information provided by the user

The Event Assistant creates multiple tables within the portal server database, and processing overhead by the Event Assistant increases the portal server CPU, disk and memory usage.

The Event Assistant is enabled by default. To reduce the processing demands for the portal server, the Event Assistant can be disabled by setting `KFW_EVENT_ASSIST=N` in the portal server configuration file (`KFWENV` on Windows, `cq.ini` on Unix or Linux). Disabling the Event Assistant causes the following behavior:

- Event records are not written to the portal server database.
- The ability for any user to attach files to an event is disabled. This is equivalent to all users not having the Attach permission.
- The My Acknowledged Events view are not updated with events that are being acknowledged.
- The Event Notes[®] view are not populated. However, to view notes and previously existing attachments, a user can still view notes through the Event Notes and Acknowledgement dialogs.
- The *Similar Events by ...* view is not populated.

KFW_EVENT_RETENTION

The number of days to keep a closed event. For example, to prune an event 2 days after it is close, specify 2. By default, no event pruning occurs. If the Event Assistant is disabled (`KFW_EVENT_ASSIST=N`), this parameter is ignored.

KFW_PRUNE_START

The time of day to start pruning data, specified in 24-hour notation (hh:mm). For example, to begin pruning data at 11:00 PM, specify 23:00. By default, no event pruning occurs. If the Event Assistant is disabled (KFW_EVENT_ASSIST=N), this parameter is ignored.

KFW_PRUNE_END

The time of day to stop pruning data, specified in 24-hour notation (hh:mm). For example, to stop pruning data at midnight, specify 24:00. By default, no event pruning occurs. If the Event Assistant is disabled (KFW_EVENT_ASSIST=N), this parameter is ignored.

You can control the size of file attachments for events either at the individual client level or at the monitoring environment level, by changing environment variables in the portal server environment file. Consider changing these if you want to restrict the number of large attachments to be held at the portal server:

KFW_ATTACHMENT_MAX

Use this parameter to control the maximum size of all the files attached to an acknowledgment. Consider editing this parameter if the size of event attachments are too large and are causing network issues, or alternatively you have excess capacity that may be used and attachments have been discarded. Enter the maximum size in bytes, such as 1000000 for 1 MB. The default value is 10000000 (10 MB). If the Event Assistant is disabled, this parameter is ignored.

KFW_ATTACHMENT_SEGMENT_MAX

This parameter enables you to set a size limitation for individual files attached to an acknowledgment. Enter the maximum size in bytes, such as 250000 for 250 KB. The default value is 1000000 (1 MB). If the Event Assistant is disabled, this parameter is ignored.

Topology update processing

The portal server maintains an internal representation of the topology of the monitoring environment, which is called the topology tree. When changes to the topology occur, the portal server updates the topology tree. Examples of topology changes that cause updates to the topology tree include:

- A new agent or monitoring server connects to a monitoring server for the first time.
- Either the hostname or the IP address for managed system changes.
- A managed system is removed or cleared from the monitoring server using either the portal client or the tacmd cleanms command.

The process of updating the topology tree can be CPU-intensive for large environments, and the portal server can cause high CPU utilization for a brief period, depending on the size of the environment and the speed of the processor running the portal server. Performance is also dependent on the number of processors or processor cores. The rebuilding of the physical navigator tree utilizes all processors on the box to improve response time. To minimize these processing requirements, the portal server batches topology updates to be processed as a group. The following portal server environment variables control how topology updates are batched together and affect the frequency of update processing for the topology tree.

KFW_CMW_UPDATE_TOPOLOGY_CUTOFF=xx

This variable sets how long the portal server waits after receiving a topology update for new updates to arrive before updating the topology tree. If a new update arrives, the portal server waits for another cutoff interval for more updates to arrive. If no new updates appear, the update processing is started for topology tree. The default cutoff value is 20 seconds. Changing the value to 60 seconds increases the possibility that topology updates are batched together and can reduce the frequency of topology tree updates.

KFW_CMW_UPDATE_TOPOLOGY_MAX_WAIT=xxx

This variable sets the maximum time the portal server waits when topology updates are arriving before updating the topology tree. If multiple topology updates arrive and are batched together across several cutoff intervals, this variable sets the maximum time the first topology update

waits before starting update processing for the topology tree. The default value is 120 seconds. Changing the value to 300 seconds allows for more batching of topology update requests.

Note: Settings made at the client level take precedence over the settings at the monitoring environment level defined here. If you have specified a setting at the client level, then any event variables defined within the portal server configuration are ignored. If you have not specified a setting at the client level, the portal server variables are used. If you have not specified any settings, the default values are used. Figure 95 shows an example of how the topology update buffering algorithm would work using the following values:

- KFW_CMW_UPDATE_TOPOLOGY_CUTOFF=120
- KFW_CMW_UPDATE_TOPOLOGY_MAX_WAIT=300




Figure 95. Topology update buffering algorithm example

The red arrows above the time line show when topology updates arrive. The blue arrows below the time line indicate when the tree rebuild would start. For the first two topology updates, there is a period of 120 seconds following their arrival when there were no additional arrivals. For the first two updates, the tree rebuild would occur after the CUTOFF interval. For the third topology update arrival, there are multiple arrivals after it, and no period of 120 seconds when no updates arrived. For these updates, the tree rebuild would start after the MAX_WAIT period of 300 seconds expired.

For machines with multiple network adapters, the recommended practice is to specify the IP address for the agent to use in the agent configuration file:

```
KDEB_INTERFACELIST=preferred_ipaddress
```

If you cannot update the agent definitions in the time required, as a short-term measure there is a Tivoli Enterprise Portal Server environment variable that you can set to prevent topology updates due to IP address changes. Use of this variable is not generally advised since it prevents the updated information from appearing in the Tivoli Enterprise Portal console.

```
KFW_CMW_DETECT_AGENT_ADDR_CHANGE=N
```

If agents are not correctly configured and generate excessive Navigator tree rebuilding, you can set this variable to have any discovery of changes or additions of IP address ignored.

For environments with more than 10 events per minute, consider disabling the Event Assistant

The Event Assistant creates multiple tables within the Tivoli Enterprise Portal Server database, and processing overhead by the Event Assistant increases the Tivoli Enterprise Portal Server CPU, disk, and memory usage. The Event Assistant is enabled by default.

For more information, see “Portal server parameter tuning” on page 475.

Perform periodic table maintenance for Tivoli Enterprise Portal Server DB2 database tables

If you are using DB2 as your Tivoli Enterprise Portal Server database, you should perform periodic table maintenance on the Tivoli Enterprise Portal Server tables, including RUNSTATS.

You can do this by turning on automatic table maintenance (AUTO_TBL_MAINT) and automatic runstats (AUTO_RUNSTATS).

Tivoli Enterprise Portal client

The Tivoli Enterprise Portal client issues requests to the Tivoli Enterprise Portal Server and renders the data that is returned. Depending on the choice of installation, the portal client can be started as a desktop application or as an applet embedded in a Web browser.

The topics in this section outline considerations for optimizing the portal client performance:

- “Tuning the portal client JVM”
- “Portal client parameter tuning” on page 479
- “Avoid displaying a large number of events in the Situation Event Console” on page 480

Tuning the portal client JVM

The memory usage of the portal client JVM increases as the size of the monitoring environment increases. If the maximum Java heap size setting is too low, the JVM spends a significant amount of time performing garbage collection. This can result in poor portal client response time and high CPU utilization.

The default value for the maximum Java heap size differs by client type, as shown in Table 78.

Table 78. Default Java heap size parameters by portal client type

Client type	Initial Java heap size (-Xms)	Maximum Java heap size (-Xmx)	Where specified
Browser	128 MB	512 MB	For portal server on Windows: <itm_install_dir>\cnb\jreconfig.js For portal server on xLinux, zLinux, or AIX: <itm_install_dir>/<arch>/cw/jreconfig.js
Desktop	128 MB	256 MB for 32-bit OS environments 512 MB for 64-bit OS environments	<ul style="list-style-type: none">• cnp.bat (Windows)• cnp.sh (Linux)• in install_dir/CNP

Table 78. Default Java heap size parameters by portal client type (continued)

Client type	Initial Java heap size (-Xms)	Maximum Java heap size (-Xmx)	Where specified
Java Web Start Desktop client	128 MB	512 MB	<itm_install_dir>\Config\tep.jnlpt Note: After modification of this file, you will need to run a re-configuration of the portal server in order for the associated Java Web Start deployment files to be regenerated.

Portal client parameter tuning

This section provides some information about parameters you might consider editing to improve portal client performance and usability.

The parameters are set in the following files according to operating system:

Windows

ITM_HOME\CNP\cnp.bat (or *cnp_instance_name.bat*)

For example: C:\IBM\ITM\CNP\cnp.bat

Linux

ITM_HOME/OS_Specific_Directory/cj/bin/cnp.sh (or *cnp_instance_name.sh*)

For example: /opt/IBM/ITM/li6243/cj/bin/cnp.sh

Be aware that on each occasion maintenance or reconfiguration takes place in your environment these files might be recreated and changes lost and need to be reapplied.

cnp.databus.pageSize

Number of rows to fetch in single logical page for any workspace table. By default 100 rows are returned. Although there is no limit to what you can set here, the larger the page size, the more memory required at the client and portal server. To increase manageability you might want to edit this value to return more rows.

cnp.attachment.total.maxsize

Use this parameter to control the maximum size of all the files attached to an acknowledgment. Consider editing this parameter if the size of event attachments are too large and are causing network issues, or alternatively you have excess capacity that might be used and attachments have been discarded. Enter the maximum size in bytes, such as 1000000 for 1 MB. The default value is 10 MB.

cnp.attachment.segment.maxsize

This parameter enables you to set a size limitation for individual files attached to an acknowledgment. Enter the maximum size in bytes, such as 250000 for 250 KB. The default value is 1 MB.

The following two parameters control the behavior when expanding items in the Navigator view which have a large number of items in the expanded list. The amount of time required to expand Navigator branches depends on how many items are in the expanded list. These parameters allow you to control how many items are expanded at a time.

cnp.navigator.branch.pagesize

The number of items to fetch on a Navigator branch expansion request. The default value is 25.

cnp.navigator.branch.threshold

The warning threshold for Navigator branch expansion requests. The default value is 100.

The following four parameters control the behavior of the Situation Event Console view when it is initially displayed on a workspace, and when it adds events to the display. The default values allow the portal client user to make selections while events are processed in the background.

cnp.siteeventcon.initial_batchsize

Maximum number of events that the situation event console will process in the first event batch cycle. The default value is 100.

cnp.siteeventcon.initial_dispatchrate

Number of milliseconds that will elapse after the first event batch cycle is processed by the situation event console. The default value is 5000 milliseconds (5 seconds).

cnp.siteeventcon.batchsize

Maximum number of events that the situation event console will process in all subsequent event batch cycles. The default value is 100.

cnp.siteeventcon.dispatchrate

Number of milliseconds that will elapse after each subsequent event batch cycle is processed by the situation event console. The default value is 1000 milliseconds (1 second).

Note: Settings made here at the client level take precedence over those at the monitoring environment level. If you have specified a setting at the client level, then any event variables defined within the portal server configuration are ignored. If you have not specified a setting at the client level, the portal server variables are used. If you have not specified any settings, the default values are used.

Avoid displaying a large number of events in the Situation Event Console

Tivoli Enterprise Portal performance starts to slow down as the number of events in the Situation Event Console view increases, especially if there are more than 1000 events.

Consider the following approaches to reduce the number of events displayed by the Tivoli Enterprise Portal:

- Review the most frequently occurring events to determine if they represent a take action item for the user.
- Consider removing the Situation Event Console view from the default workspaces if you do not use it.
- If events are forwarded and managed using Tivoli Enterprise Console or Omnibus, and if event information is not needed in the Tivoli Enterprise Portal Navigator tree and situation event consoles, you can dissociate situations which eliminates Tivoli Enterprise Portal and Tivoli Enterprise Portal Server processing related to the event.

Tivoli Data Warehouse

The Tivoli Data Warehouse is the most intensely used component of the Tivoli Monitoring infrastructure. The warehouse must support a large volume of data transactions during every warehousing period, daily summarization and pruning of the warehoused data and multiple queries that often return result sets that are thousands of rows long. The processing power required to support the warehouse database and the Summarization and Pruning Agent can be significant.

This section provides guidelines for setting your Tivoli Data Warehouse. If you are interested in additional information, see the Tivoli Management Services Warehouse and Reporting Redbook.

The following topics outline processes for optimizing the performance of the warehousing process:

- “Historical data collection” on page 481

- “Warehouse Proxy Agent” on page 482
- “Summarization and Pruning Agent” on page 486
- “Database tuning” on page 489

For all but the simplest environments, database administrator skills are needed to manage the warehouse database

Database tuning is a complex task, and for most databases it requires the skills of a database administrator. For the Tivoli Data Warehouse, a database located on a single disk using default setting database parameters is only suitable for small test environments.

For all other environments, careful planning, monitoring, and tuning are required to achieve satisfactory performance.

Historical data collection

This section contains recommendations related to historical data collection and controlling the volume of data that is generated and loaded into the Tivoli Data Warehouse.

- Do not start collecting historical data for an attribute group without first estimating the data volume that will be generated and the disk space requirements on both the agent and the warehouse database. The Warehouse load projections spreadsheet helps you make these calculations. You can find this spreadsheet by searching for “warehouse load projections” or the navigation code “1TW10TM1Y” in IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>).
- Collect historical data only for attribute groups where there is a planned use for the information. Do not collect data that will not be used.
- The number of rows per day generated by a monitoring agent collecting historical data for an attribute group can be calculated with the following formula:

$$(60 / \text{collection interval}) * 24 * (\# \text{ instances at each interval})$$

Where:

60 Represents the 60 minutes in an hour.

collection interval The data collection interval, in minutes. This value can be 1, 5, 15, 30, 60, or 1440 (1 day).

24 Represents 24 hours in one day.

instances at each interval The number of instances recorded at each interval.

The two variables in this formula are the collection interval and the number of instances recorded at each interval.

- The collection interval is a configuration parameter specified in the Historical Data Collection configuration dialog.
- The number of instances recorded at each interval depends on the nature of the attribute group and the managed system being monitored. Some attribute groups, such as NT_Memory, generate a single row of data per collection interval. Most attribute groups, however, generate multiple rows of data, one row for each monitoring instance (for example, one per CPU, one per disk, and so on). Certain attribute groups can be *instance-intense*, generating dozens or hundreds of rows of data per collection interval. Examples of instance-intense attribute groups would be those reporting information at the process level, thread level, disk level (for file servers), or network connection level.

For attribute groups that return multiple rows, the number of instances recorded at each interval is configuration-dependent, and can be different from one monitoring environment to another. The Warehouse load projections spreadsheet requires the user to specify the number of instances recorded at each interval. There are several approaches that you can use to come up with this number.

- Using the portal client, build a table view for the monitoring agent and define a query to obtain data from the required attribute group. The number of rows shown in the table view is the number

of instances that would be recorded at each interval. For details on how to define table views in the portal client, see the *IBM Tivoli Monitoring Tivoli Enterprise Portal User's Guide*.

- Issue a SOAP call to collect data for this attribute group from the monitoring agent. The number of data rows returned by the SOAP call is the number of instances that would be recorded at each interval. For details on how to issue SOAP calls, see Appendix A “Tivoli Enterprise Monitoring Web services” in the *IBM Tivoli Monitoring Administrator's Guide*.
- If you have a test environment, you can create a monitoring server to use in historical data collection testing. Enable historical data collection for the required attribute group under this remote monitoring server, and configure a representative agent to connect to this monitoring server. When the agent uploads data to the Warehouse Proxy Agent, you can query the WAREHOUSELOG table to see how many rows were written by the agent for the attribute group.

To minimize the data volume (rows per day) generated by a monitoring agent for an attribute group, consider the following two recommendations:

- Use the longest data collection interval (1, 5, 15, 30, 60 or 1440 minutes) that will provide the required level of information.
- Avoid or minimize data collection for instance-intense attribute groups, which can generate many rows per data collection interval.
- When configuring historical data collection for an attribute group, you specify which monitoring servers collect the data. Since historical data collection is configured at the monitoring server level, restrict the number of agents collecting data for an attribute group, if possible, by enabling historical data collection on a subset of the monitoring servers.
- To minimize the performance impact on the monitoring server, configure historical data to keep short-term history files at the agent, if possible, rather than at the monitoring server.
- Enable warehouse collection only for attribute groups where there is a planned use for the information. For attribute groups with historical data collection enabled but not configured for warehouse collection, you must schedule regular tasks to prune the short-term history files (the supplied programs are described in the *IBM Tivoli Monitoring Administrator's Guide*).
- To spread the warehouse collection load across the day, configure warehouse collection to occur hourly rather than daily.

Warehouse Proxy Agent

The Warehouse Proxy Agent provides a means of translating and transferring historical data collected by other monitoring agents to the Tivoli Data Warehouse.

The amount of historical data generated can be huge, particularly in environments with thousands of monitoring agents or when instance intense attribute groups are enabled.

This section describes Warehouse Proxy Agent configuration parameters that affect performance and throughput. For most environments, the default settings for these parameters provide good performance. If the Warehouse Proxy Agent starts to encounter errors when exporting data, these parameters may need adjustment. If the Warehouse Proxy Agent is not able to export data into the warehouse, or if it cannot keep up with the volume of export requests, export errors will be encountered at the storage location. The storage location can be the monitoring agent or the monitoring server, whichever is specified in the historical data collection configuration dialog.

Configuration parameters for the Warehouse Proxy Agent are set in the following files according to operating system:

Windows

`ITM_HOME\TMAITM6\khdenv`

For example: `C:\IBM\ITM\TMAITM6\khdenv`

Linux

UNIX

`ITM_HOME/config/hd.ini`

For example: /opt/IBM/ITM/config/hd.ini

Be aware that whenever maintenance or reconfiguration takes place in your environment, these files might be recreated and configuration changes can be lost and require reapplication.

Warehouse Proxy internals

To understand the performance-related configuration parameters, it is first helpful to understand how the Warehouse Proxy Agent collects and transfers data to the warehouse. The Warehouse Proxy Agent comprises three internal components; the intelligent remote agent (IRA) communication framework, the work queue, and the exporter threads.

These three components work together to collect, translate, and transmit historical data to the warehouse. Historical data flows from one component to the next, undergoing specific processing at each step before being passed on.

Tuning the Warehouse Proxy Agent on AIX and Linux systems

For a Warehouse Proxy Agent running on AIX or Linux, make sure that the user limit value for number of open file descriptors (the `nfiles` parameter) is set higher than the number of agents that are uploading data through the Warehouse Proxy Agent.

See “File descriptor (maxfiles) limit on UNIX and Linux systems” on page 137 for more details.

Warehouse Proxy Agents running on AIX and Linux systems use a Java virtual machine (JVM), which uses the JDBC (Java Database Connectivity) interface to communicate with the warehouse database. If the maximum Java heap size of the JVM is set to a low value, performance can be degraded by frequent garbage collections.

The maximum Java heap size is controlled by the `-Xmx` option. By default, this option is not specified in the Warehouse Proxy Agent configuration file. If this option is not specified, the default value used by Java applies as follows:

- AIX** Half the available memory with a minimum of 16 MB and a maximum of 512 MB.
- Linux** Half the available memory with a minimum of 16 MB and a maximum of 512 MB.

Note: The above values are from the *IBM Developer Kit and Runtime Environment, Java Technology Edition, Version 6 Diagnostics Guide*.

To set the size of maximum Java heap size for the Warehouse Proxy Agent, edit the `hd.ini` configuration file and modify the `KHD_JAVA_ARGS` variable as shown below:

```
KHD_JAVA_ARGS=-Xmx256m
```

A maximum Java heap size of 256 megabytes is more than adequate for most environments.

Setting `KHD_JNI_VERBOSE=Y` in the configuration file will enable logging of the garbage collector's actions. If the Java log contains an excessive number of garbage collection entries during a single warehousing interval, consider increasing the size of the Java heap.

Using multiple Warehouse Proxy Agents

Tivoli Monitoring supports multiple warehouse proxies within a single hub monitoring server environment. The provision for multiple warehouse proxies provides for greater scalability and performance in historical data collection, and more importantly, improves reliability by providing a failover mechanism.

If a Warehouse Proxy is unavailable, data can be inserted into the warehouse by a different Warehouse Proxy Agent (if the agents are configured properly for failover). “Installing and configuring multiple Warehouse Proxy Agents” on page 656 contains detailed information about configuring multiple Warehouse Proxy Agents.

- If you are collecting and warehousing historical data in a monitoring environment with more than 1500 monitoring agents, consider using multiple Warehouse Proxy Agents to handle the volume of data that will be uploaded and inserted into the warehouse.
- To reduce the number of servers running Tivoli Monitoring components, you can install the Warehouse Proxy Agent on the same servers running the remote monitoring servers. Servers running both a monitoring server and a Warehouse Proxy Agent should have two processors.
- By default, each Warehouse Proxy Agent opens 10 exporter threads and 10 database connections for inserting data into the warehouse database (controlled by the **KHD_EXPORT_THREADS** and **KHD_CNX_POOL_SIZE** parameters, respectively). If you are using more than five Warehouse Proxy Agents, consider reducing the number of exporter threads and database connections for each Warehouse Proxy Agent, to limit the number of database connections inserting data into the warehouse database.

Check ulimit settings for open file descriptors (UNIX/Linux only)

The Warehouse Proxy Agent is similar to the Tivoli Enterprise Monitoring Server in that it can use a large number of file descriptors, especially in a large environment.

The same recommendation for ulimit settings for the Tivoli Enterprise Monitoring Server applies to the Warehouse Proxy Agent as well. For more information, see “Check ulimit settings for open file descriptors (UNIX/Linux only)” on page 471.

Disable logging to the WAREHOUSELOG table and use trace to monitor Warehouse Proxy Agent export activity

During normal processing, the WAREHOUSELOG table is appended each time an agent successfully exports a batch of attribute group rows to the warehouse database. The entries in the WAREHOUSELOG table record the agent name, attribute group name, number of rows exported, the Warehouse Proxy Agent involved, and the time of the export.

For large environments, the WAREHOUSELOG table can grow very large, and managing the table can add unnecessary overhead for the Summarization and Pruning Agent. A more efficient approach for monitoring Warehouse Proxy Agent export activity is to disable logging to the WAREHOUSELOG table, and instead turn on a trace setting that records similar information to the Warehouse Proxy Agent trace log.

Note: Warehouse Proxy Agent logging to the WAREHOUSELOG table is disabled by default. The self-monitoring workspaces provided by the Warehouse Proxy Agent provide sufficient information to determine if the agent is operating correctly.

To disable logging to the WAREHOUSELOG table manually, add the following to the Warehouse Proxy Agent configuration file:

```
KHD_WHLOG_ENABLE=N
```

To turn on tracing for Warehouse Proxy Agent export activity, set the trace options on the KBB_RAS1 entry in the Warehouse Proxy Agent configuration file:

```
KBB_RAS1=ERROR (UNIT:khdxdbex OUTPUT)
```

This trace setting writes an entry to the trace log every time an agent exports a batch of attribute group rows to the warehouse database. The trace entry contains similar information to the information that is written to the WAREHOUSELOG, that is agent name, attribute group name, and number of rows.

You can tail the Warehouse Proxy Agent trace log to monitor Warehouse Proxy Agent export activity in real-time.

Example:

```
# tail -20f prfr1s33_hd_4c582197-02.log
(4C5DC26D.00EA-14:khdxdbex.cpp,3022,"endProcessSample") Inserted 4 rows of data into
"Linux_VM_Stats" (LNxVM appl KLZ) for "177:prfr2s34:LZ" , status 0
```

```
(4C5DC26D.00EB-BD:khdxdbex.cpp,3022,"endProcessSample") Inserted 4 rows of data into
"Linux_VM_Stats" (LNXVM appl KLZ) for "228:prfr2s34:LZ" , status 0
(4C5DC26D.00EC-16:khdxdbex.cpp,3022,"endProcessSample") Inserted 16 rows of data into
"Linux_Network" (LNXNET appl KLZ) for "125:prfr2s34:LZ" , status 0
```

The trace setting described in this section writes export information in the trace log, which does wrap and will not grow indefinitely.

If using DB2 for z/OS for Tivoli Data Warehouse database, increase the idle thread timeout parameter

When using DB2 for z/OS 9.1 for the Tivoli Data Warehouse database, the Warehouse Proxy Agent can encounter repeated disconnections from the database. The default idle thread timeout value (DSN6FAC IDTHTOIN in DSNZPxxx) is 120 seconds.

The Warehouse Proxy Agent uses a pool of database connections to process export requests from monitoring agents. The warehousing interval used by agents can be set to values ranging from 15 minutes up to 24 hours. The database connections are idle between export requests, and if the idle thread timeout value is less than the warehousing interval, the database connections might timeout. This results in numerous error messages written to the Warehouse Proxy Agent log. The Warehouse Proxy Agent *Statistics* workspace also shows many disconnections in the *Failure / Disconnections* view.

To avoid repeated disconnections, consider increasing the DB2 idle thread timeout value to a value higher than the warehousing interval. Specifying a value of 0 disables time-out processing. If time-out processing is disabled, idle server threads remain in the system and continue to hold their resources, if any.

Configuration parameter changes used in large-scale environments

Configuration parameters for the Warehouse Proxy Agent are set in the following files according to operating system.

Windows

`ITM_HOME\TMAITM6\khdenv`

For example: `C:\IBM\ITM\TMAITM6\khdenv`

Linux

UNIX

`ITM_HOME/config/hd.ini`

For example: `/opt/IBM/ITM/config/hd.ini`

This section contains configuration settings to the Warehouse Proxy Agent configuration files, that you might consider to improve performance in large-scale environments.

KHD_WAREHOUSE_TEMS_LIST=REMOTE_<hostname>

This value is needed when using multiple Warehouse Proxy Agents in a configuration. This value is set so that the Warehouse Proxy Agent services agents that are connected to the local remote monitoring server.

KHD_EXPORT_THREADS=5

With multiple Warehouse Proxy Agents in the configuration, you can lower this value from the default of 10 to reduce the number of concurrent connections to the warehouse database.

KHD_CNX_POOL_SIZE=5

With multiple Warehouse Proxy Agents in the configuration, you can lower this value from the default of 10 to reduce the number of concurrent connections to the warehouse database.

KBB_RAS1=ERROR (UNIT:khdxdbex OUTPUT)

You can add the trace setting (*UNIT:khdxdbex OUTPUT*) to the KBB_RAS1 entry. This setting causes export messages to write to the Warehouse Proxy Agent trace logs, which provides useful information for performance analysis with negligible overhead. This setting is useful if you are specifying `KHD_WHLOG_ENABLE=N`.

DB2CODEPAGE=1208 (DB2 only)

This setting avoids an initialization error message of *DB2 client code page not found*, which appears in the Warehouse Proxy workspace. The value that you use for the DB2CODEPAGE variable must be the same as what you specify on your CREATE DATABASE statement. For more information about setting the DB2CODEPAGE variable, see the IBM DB2 Database information center at <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp>.

NLS_LANG=AMERICAN_AMERICA.AL32UTF8 (Oracle only)

This setting is used in place of DB2CODEPAGE if using Oracle for the warehouse database.

KHD_JAVA_ARGS=-Xmx256m (UNIX/Linux only)

Sets the maximum Java heap size for the khdxprtj Java virtual machine. This setting is useful in lowering the process size (SZ) and virtual size (VSZ) for the khdxprtj process on the Tivoli Enterprise Monitoring Server.

Summarization and Pruning Agent

The Summarization and Pruning Agent is a multi-threaded, Java based application. It interacts with the warehouse using a JDBC driver appropriate for the database of the warehouse. The number of worker threads available and the heap size of the JVM affect the performance of the Summarization and Pruning Agent and the length of time of its processing runs.

The installation location of the Summarization and Pruning Agent is another important aspect of Summarization and Pruning Agent performance tuning.

On Windows, the Summarization and Pruning Agent environment parameters are set in the following configuration files according to operating system:

Windows

`ITM_HOME\TMAITM6\KSYENV`

For example: `C:\IBM\ITM\TMAITM6\KSYENV`

Linux

UNIX

`ITM_HOME/config/sy.ini`

For example: `/opt/IBM/ITM/config/sy.ini`

Whenever maintenance or reconfiguration takes place in your environment, these files might be recreated, and configuration changes can be lost and require reapplication.

Number of worker threads

The Summarization and Pruning Agent creates a pool of worker threads during initialization for performing the summarization and pruning tasks. Each worker thread operates independently of all others, concentrating on one attribute group and its associated warehouse tables. After a worker thread finishes an attribute group, it locates the next attribute group scheduled for processing. If there are no more attribute groups to process, the thread ends, leaving the remaining threads to finish their work.

This threading model provides concurrency while minimizing lock activity, since threads are working on different tables. If there are a few attribute group tables that are considerably larger than the other attribute group tables, the run time of the Summarization and Pruning Agent will be gated by the processing times for the largest attribute group tables.

The number of worker threads can be set in the **Additional Parameters** tab of the configuration window, or by setting the variable KSY_MAX_WORKER_THREADS in the configuration file (KSYENV on Windows, sy.ini on UNIX or Linux).

- The default value is 2.

- The suggested number of worker threads is 2 or 4 times the number of processors on the host server. If the Summarization and Pruning Agent is running on a separate server from the warehouse database server, set the value based on the number of processors on the warehouse database server.
- Configuring more threads than attribute groups will not decrease the processing time, because each thread works on one attribute group at a time.

Setting the maximum Java heap size

The Summarization and Pruning Agent runs in a Java virtual machine (JVM), which uses the JDBC (Java Database Connectivity) interface to communicate with the warehouse database. If the maximum Java heap size of the JVM is set to a low value, performance can be degraded by frequent garbage collections.

The maximum Java heap size is controlled by the **-Xmx** option. By default, this option is not specified in the Summarization and Pruning Agent configuration file. If this option is not specified, the default value used by Java applies as follows:

- **AIX** Half the available memory with a minimum of 16 MB and a maximum of 512 MB.
- **Linux** Half the available memory with a minimum of 16 MB and a maximum of 512 MB.
- **Windows** Half the real memory with a minimum of 16 MB and a maximum of 2 GB.

Note: The above values are from the *IBM Developer Kit and Runtime Environment, Java Technology Edition, Version 6 Diagnostics Guide*.

To set the size of maximum Java heap size for the Summarization and Pruning Agent, edit the configuration file (KSYENV on Windows, sy.ini on UNIX or Linux) and modify the KSZ_JAVA_ARGS variable as shown below:

```
KSZ_JAVA_ARGS=-Xmx512m
```

- A maximum Java heap size of 512 megabytes (shown in the example above) is adequate for most environments.
- In addition to the “**mx**” Java parameter, you can also specify the **-verbose:gc** Java runtime parameter, which causes diagnostic messages related to garbage collection to be written to the log. If there are an excessive number of garbage collection entries consider increasing the size of the Java heap.
- For more information about Java heap tuning parameters, see the *IBM Developer Kit and Runtime Environment, Java Technology Edition, Version 6 Diagnostics Guide*, which is available from <http://www.ibm.com/developerworks/java/jdk/diagnosis>.

Enabling more detailed trace in log files

The Summarization and Pruning Agent Java internal trace files contain diagnostic messages showing the number of rows read and pruned for each attribute group and agent. The names of these files are of the form `hostname_sy_java_timestamp-n.log`. For large environments, these trace files might be overwritten during a single Summarization and Pruning Agent processing cycle, and some of the diagnostic information may be lost.

By default, the Java-based internal trace wraps at 5 files, and each file contains 300000 lines. To change the default values, you can specify Java runtime parameters in the Summarization and Pruning Agent configuration file (KSYENV on Windows, sy.ini on UNIX or Linux):

```
KSZ_JAVA_ARGS=-Dibm.tdw.maxNumberDetailTraceFiles=A
-Dibm.tdw.maxLinesForDetailTraceFile=B
```

where:

- A** Specifies the maximum number of Java-based internal trace files that can exist at any one time for a single launch.
- B** Specifies the maximum number of lines per Java-based internal trace file.

Consider increasing these log parameters so that you have a few days worth of data in the logs for diagnostic purposes.

Consider disabling shifts and vacations

The Summarization and Pruning Agent configuration settings can be specified in the configuration window, which is described in the various Summarization and Pruning Agent installation and configuration steps.

See Chapter 27, “Tivoli Data Warehouse solutions: common procedures,” on page 643 for more information. The **Work Days** tab in this configuration window allows you to specify shift information and vacation settings.

When you enable and configure shifts, IBM Tivoli Monitoring produces three separate summarization reports:

- Summarization for peak shift hours
- Summarization for off-peak shift hours
- Summarization for all hours (peak and off-peak)

Similarly, when you enable and configure vacations, IBM Tivoli Monitoring produces three separate summarization reports:

- Summarization for vacation days
- Summarization for non-vacation days
- Summarization for all days (vacation and non-vacation)

Enabling shifts and vacations causes increased processing and database space usage by the Summarization and Pruning Agent. If you do not require the separate shift and vacation summarization reports, verify that shift information and vacation settings are not enabled in the Summarization and Pruning Agent configuration window. Alternatively, to disable shifts and vacations, you can specify the following parameter settings in the configuration file (KSYENV on Windows, sy.ini on UNIX or Linux):

KSY_SHIFTS_ENABLED=N

KSY_VACATIONS_ENABLED=N

Minimize the number of summarization intervals used

Keep the number of summarization intervals used, such as hourly, daily, weekly, to a minimum since each additional intervals increase the Summarization and Pruning Agent workload significantly.

The summarization intervals are processed independently, so to use daily summarization it is not necessary to also use hourly summarization.

Set the max worker threads and max rows per transaction appropriately for your environment

In the large scale environment, you might want to change the parameters in the Summarization and Pruning Agent configuration file (sy.ini for UNIX/Linux or ksyenv for Windows).

KSY_MAX_WORKER_THREADS

This parameter specifies the number of worker threads for processing attribute group tables. The default value is 2. For a large database server with multiple processors, increasing this value might improve concurrency and throughput. A good starting value is two times the number of processors on the database server.

KSY_MAX_ROWS_PER_TRANSACTION

This parameter specifies the maximum number of rows per database transaction, and can be thought of as a batching factor. When the number of worker threads is greater than 1, the batching factor for a given worker thread is equal to the maximum rows per transaction divided by the number of worker threads. When you set the number of worker threads, you should also consider increasing the value, to maintain a batching factor of 200 or more.

Database tuning

Database tuning is a complex task, and for important databases, it requires the skills of a database administrator. For the Tivoli Data Warehouse, a database located on a single disk using default database parameters is suitable only for small test environments. For all other environments, careful planning, monitoring, and tuning are required to achieve satisfactory performance.

There are a number of sources of database configuration and tuning information that should be helpful in the planning, monitoring, and tuning of the Tivoli Data Warehouse:

1. "Understanding the disk requirements for your database" on page 513 describes factors to consider in planning the disk subsystem to support the Tivoli Data Warehouse.
2. The paper "Relational Database Design and Performance Tuning for DB2 Database Servers" is available from the Tivoli Integrated Service Management Library by searching for "database tuning" or navigation code "1TW10EC02" at IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>). This paper is a concise reference describing some of the major factors that affect DB2 performance. This document is a good starting point to use before referencing more detailed information in manuals and Redbooks® devoted to DB2 performance. While DB2 specific, many of the concepts are applicable to Relational Databases in general, such as Oracle and Microsoft SQL Server.
3. The "Tivoli Data Warehouse tuning" chapter of the Redbook Tivoli Management Services Warehouse and Reporting (SG24-7443) builds upon the previously referenced Tivoli Integrated Service Management Library paper, supplementing it with information about Oracle and Microsoft SQL Server. This chapter is almost 100 pages in length.
4. The "Optimizing the performance" chapter of the Redbook IBM Tivoli Monitoring Implementation and Performance Optimization for Large Scale Environments (SG24-7443) contains a section on database tuning considerations for the Tivoli Data Warehouse. This section makes suggestions about specific tuning parameters for DB2, Oracle and MS SQL. At approximately 12 pages, this section is much shorter than the tuning chapter in the previously referenced Redbook (item number 3).

The remainder of this section is a summarized version of material from the Tivoli Integrated Service Management Library paper in number 2 above, which has been supplemented by identifying specific parameters relevant to Tivoli Data Warehouse and providing some suggested ranges of values.

Relational database design and performance tuning for DB2 Database servers

This section explains some of the major factors that affect DB2 performance on distributed platforms such as Windows, UNIX, and Linux. Many of the concepts are applicable to the topic of relational databases such as Oracle and Microsoft SQL Server.

This information does not replace detailed information about DB2 performance that is available in various manuals and Redbooks. You can find the Redbooks at <http://www.redbooks.ibm.com>. You can find the DB2 library of documents at <http://pic.dhe.ibm.com/infocenter/db2luw/v9r7/index.jsp>.

Terminology

The terms in this section are useful for understanding performance issues.

Throughput

The amount of data transferred from one place to another or processed in a specified amount of time. Data transfer rates for disk drives and networks are measured in terms of throughput. Typically, throughputs are measured in kilobytes per second, Mbps, and Gbps.

Optimizer

When an SQL statement is run, the SQL compiler must determine the access plan to the database tables. The optimizer creates this access plan, using information about the distribution of data in specific columns of tables and indexes if these columns are used to select rows or join tables. The

optimizer uses this information to estimate the costs of alternative access plans for each query. Statistical information about the size of the database tables and available indexes heavily influences the optimizer estimates.

Clustered Index

An index whose sequence of key values closely corresponds to the sequence of rows that are stored in a table. Statistics that the optimizer uses measure the degree of this correspondence.

Cardinality

The number of rows in the table or for indexed columns the number of distinct values of that column in a table.

Prefetch

An operation in which data is read before its use when its use is anticipated. DB2 supports the following mechanisms:

Sequential prefetch

A mechanism that reads consecutive pages into the buffer pool before the application requires the pages.

List prefetch or list sequential prefetch

Prefetches a set of non-consecutive data pages efficiently.

Performance factors

The performance factors in this section, which are thoroughly detailed in subsequent sections, affect overall performance of any application.

- Database Design
- Application Design
- Hardware Design and Operating System Usage

This section identifies areas where you can influence performance of the Tivoli Data Warehouse database.

Database design details

Key factors for database design include table spaces, buffer pools, and logging.

The subsequent sections include information about the following topics:

- Files that are created to support and manage your database
- Amount of required space for storing your data
- Determining how you use the table spaces that are required to store your data
- Setting the various database parameters to fit your environment

Table spaces:

A *table space* is a physical storage object that provides a level of indirection between a database and the tables stored within the database. It is made up of a collection of containers into which database objects are stored.

A *container* is an allocation of space to a table space. Depending on the table space type, the container can be a directory, device, or file. The data, index, long field, and LOB portions of a table can be stored in the same table space, or can be individually broken out into separate table spaces.

When you are working with database systems, the main objective is your ability to store and retrieve data as quickly and efficiently as possible. One important consideration when designing your database or analyzing a performance problem on an existing database is the physical layout of the database itself.

DB2 provides support for two types of table spaces:

System Managed Space (SMS)

Stores data in operating system files. This type is an excellent choice for general-purpose use, providing good performance with little administration cost.

Database Managed Space (DMS)

Includes database manager control of the storage space. A list of devices or files is selected to belong to a table space when it is defined. The DB2 database manager manages the space on those devices or files. Some additional administration cost is incurred with DMS table spaces because the size of the pre-allocated files is monitored and adjusted. Altering an existing container or adding a new container to it can easily increase the DMS table space size.

Performance and table space types:

DMS table spaces usually perform better than SMS table spaces because they are pre-allocated and do not use up time extending files when new rows are added. DMS table spaces can be either raw devices or file system files. DMS table spaces in raw device containers provide the best performance because double buffering does not occur. *Double buffering*, which occurs when data is buffered first at the database manager level and subsequently at the file system level, might be an additional cost for file containers or SMS table spaces.

If you use SMS table spaces, consider using the **db2empfa** command on your database. The **db2empfa** command, which runs the Enable Multipage File Allocation tool, runs multipage file allocation for a database. With multipage file allocation enabled for SMS table spaces, disk space is allocated one extent rather than one page at a time, improving INSERT throughput. In version 8 of DB2, this parameter is turned on by default.

If you are using a RAID device, create the table space with a single container for the entire array. When a database is created in DB2, a default table space called USERSPACE1 is created, and by default, Tivoli Monitoring uses this table space when creating its tables in the DB2 database. You can create a new default table space in DB2 by creating a table space with the name IBMDEFAULTGROUP. If a table space with that name, and a sufficient page size, exists when a table is created without the IN table space clause, it is used. You can create the table space in a different location, for example, a RAID array. You could also create it as a DMS table space if you prefer, as in the following example:

```
CREATE REGULAR TABLESPACE IBMDEFAULTGROUP IN DATABASE
PARTITION GROUP IBMDEFAULTGROUP PAGESIZE 4096 MANAGED
BY SYSTEM
  USING ('E:\DB2\NODE0000\SQL00001\IBMDEFAULTGROUP')
  EXTENTSIZE 32
  PREFETCHSIZE AUTOMATIC
  BUFFERPOOL IBMDEFAULTBP
  OVERHEAD 12.670000
  TRANSFERRATE 0.180000
  DROPPED TABLE RECOVERY ON;
```

File system caching on a Windows system:

For Windows systems caching, the operating system might cache pages in the file system cache for DMS file containers and all SMS containers. For DMS device container table spaces, the operating system does not cache pages in the file system cache.

On Windows, the DB2 registry variable DB2NTNOCACHE specifies whether DB2 opens database files with the NOCACHE option. If DB2NTNOCACHE is set to ON, file system caching is eliminated. If DB2NTNOCACHE is set to OFF, the operating system caches DB2 files. This standard applies to all data except for files that contain LONG FIELDS or LOBs. Eliminating system caching allows more available memory for the database, and the buffer pool or SORTHEAP can be increased.

Buffer pools:

A *buffer pool* is an area of memory into which database pages are read, modified, and held during processing.

On any system, accessing memory is faster than disk I/O. DB2 uses database buffer pools to attempt to minimize disk I/O. Although the amount of memory to dedicate to the buffer pool varies, in general it is true that more memory is preferable. Start with 50 - 75% of your system's main memory devoted to buffer pools if the machine is a dedicated database server. Because it is a memory resource, buffer pool usage must be considered along with all other applications and processes that are running on a server. If your table spaces have multiple page sizes, create only one buffer pool for each page size.

Logging:

Maintaining the integrity of your data is important. All databases maintain log files that record database changes. DB2 *logging* involves a set of primary and secondary log files that contain log records that show all changes to a database.

The database log is used to roll back changes for units of work that are not committed and to recover a database to a consistent state.

DB2 provides the following strategies for logging:

- Circular logging
- Log retention logging

Circular logging:

Circular logging is the default log mode in which log records fill the log files and later overwrite the initial log records in the initial log file.

The overwritten log records are not recoverable. This type of logging is not suited for a production application.

Log retain logging:

With log retain logging, each log file is archived when it fills with log records.

New log files are made available for log records. Retaining log files enables roll-forward recovery, which reapplies changes to the database based on completed units of work (transactions) that are recorded in the log. You can specify that roll-forward recovery is done to the end of the logs, or to a specific point in time before the end of the logs. Because DB2 never directly deletes archived log files, the application is responsible for maintaining them (for example, archiving, purging, and so on).

Log performance:

Ignoring the logging performance of your database can be a costly mistake, especially in terms of time. Optimize the placement of the log files for write and read performance, because the database manager must read the log files during database recovery.

To improve logging performance, use the following suggestions:

- Use the fastest disks available for your log files. Use a separate array or channel if possible.
- Use Log Retain logging.
- Mirror your log files.

Increase the size of the database configuration Log Buffer parameter (LOGBUFSZ). This parameter specifies the amount of the database heap to use as a buffer for log records before writing these records to disk. The log records are written to disk when one of the following points occurs:

- A transaction commits, or a group of transactions commit, according to the definition in the MINCOMMIT configuration parameter.
- The log buffer is full.
- Another internal database manager event occurs.
- Buffering the log records supports more efficient logging file I/O because the log records are written to disk less frequently and more log records are written each time.
- Tune the Log File Size (LOGFILSIZ) database configuration parameter so that you are not creating excessive log files.

Database maintenance:

A regularly scheduled maintenance plan is essential for maintaining peak performance of your system.

Regular maintenance, which involves running the REORG, RUNSTATS, and REBIND facilities in that order on the database tables, is a critical factor in the performance of a database environment. Implement at least a minimum weekly maintenance schedule.

REORG:

After many INSERT, DELETE, and UPDATE changes to table data, often involving variable length columns activity, the logically sequential data might be on non-sequential physical data pages. The database manager must perform additional read operations to access data. You can use the **REORG** command to reorganize DB2 tables, eliminating fragmentation and reclaiming space.

Regularly scheduled REORGs improve I/O and significantly reduce elapsed time. Implement a regularly scheduled maintenance plan.

DB2 provides the following types of REORG operation, classic REORG and In-Place REORG. If you have an established database maintenance window, use the classic REORG. If you operate a 24 by 7 operation, use the In-Place REORG.

- Classic REORG
 - The fastest method of REORG
 - Indexes rebuilt during the reorganization
 - Ensures perfectly ordered data
 - Access limited to read-only during the UNLOAD phase, with no access during other phases
 - Not restartable
- In-Place REORG
 - Slower than the Classic REORG and takes more time to complete
 - Does not ensure perfectly ordered data or indexes
 - Requires more log space
 - Can be paused and restart
 - Can allow applications to access the database during reorganization

RUNSTATS:

The DB2 optimizer uses information and statistics in the DB2 catalog to determine optimal access to the database based on the provided query. Statistical information is collected for specific tables and indexes in the local database when you run the RUNSTATS utility.

When significant numbers of table rows are added or removed, or if data in columns for which you collect statistics is updated, execute **RUNSTATS** again to update the statistics.

Use the **RUNSTATS** utility to collect statistics in the following situations:

- When data was loaded into a table and the appropriate indexes were created
- When you create a new index on a table. Execute **RUNSTATS** for the new index only if the table was not modified since you last ran **RUNSTATS** on it.
- When a table has been reorganized with the REORG utility
- When the table and its indexes have been extensively updated by data modifications, deletions, and insertions. "Extensive" in this case might mean that 10 to 20 percent of the table and index data was affected.
- Before binding or rebinding, application programs whose performance is critical
- When you want to compare current and previous statistics. If you update statistics at regular intervals, you can discover performance problems early.
- When the prefetch quantity is changed
- When you have used the REDISTRIBUTE DATABASE PARTITION GROUP utility

The **RUNSTATS** command has several formats that primarily determine the depth and breadth of statistics that are collected. If you collect more statistics, the command takes more time to run. The following options are included:

- Collecting either SAMPLED or DETAILED index statistics
- Collecting statistics on all columns or only columns used in JOIN operations
- Collecting distribution statistics on all, key, or no columns. Distribution statistics are useful when you have an uneven distribution of data on key columns.

Take care when running **RUNSTATS**, because the collected information impacts the selection of access paths by the optimizer. Implement **RUNSTATS** as part of a regularly scheduled maintenance plan if some of the conditions occur. To ensure that the index statistics are synchronized with the table, execute **RUNSTATS** to collect table and index statistics at the same time.

Consider some of the following factors when deciding what type of statistics to collect:

- Collect statistics only for the columns that join tables or in the WHERE, GROUP BY, and similar clauses of queries. If these columns are indexed, you can specify the columns with the ONLY ON KEY COLUMNS clause for the **RUNSTATS** command.
- Customize the values for num_freqvalues and num_quantiles for specific tables and specific columns in tables.
- Collect detailed index statistics with the SAMPLE DETAILED clause to reduce the amount of background calculation performed for detailed index statistics. The SAMPLE DETAILED clause reduces the time required to collect statistics and produces adequate precision in most cases.
- When you create an index for a populated table, add the COLLECT STATISTICS clause to create statistics as the index is created.

REBIND:

After running **RUNSTATS** on your database tables, you must rebind your applications to take advantage of those new statistics.

Rebinding ensures that DB2 is using the best access plan for your SQL statements. Perform a **REBIND** after running **RUNSTATS** as part of your normal database maintenance procedures. The type of SQL that you are running determines how the rebind occurs.

DB2 provides support for the following types of SQL:

- Dynamic SQL

- Static SQL

Dynamic SQL:

Dynamic SQL statements are prepared and executed at run time. In dynamic SQL, the SQL statement is contained as a character string in a host variable and is not precompiled. Dynamic SQL statements and packages can be stored in one of the DB2 caches.

A rebind occurs under the following conditions when you are using dynamic SQL:

- If the statement is not in the cache, the SQL optimizer 'binds' the statement and generates a new access plan.
- If the statement is in the cache, no rebind occurs. To clear the contents of the SQL cache, use the `FLUSH PACKAGE CACHE SQL` statement.

Static SQL:

Static SQL statements are embedded within a program, and are prepared during the program preparation process before the program is executed. After preparation, a static SQL statement does not change, although values of host variables that the statement specifies can change. These static statements are stored in a DB2 object called a *package*.

A rebind occurs under the following conditions when you are using static SQL:

- Explicitly, if an explicit `REBIND` package occurs
- Implicitly, if the package is marked 'invalid', which can happen if an index that the package was using was dropped.

Application design details

The Tivoli Monitoring application was designed and tested with performance and scalability in mind. Consider reuse of the database connection using the Tivoli Monitoring connection pooling feature.

Connection pooling is a process in which DB2 drops the inbound connection with an application that requests disconnection, but keeps the outbound connection to the host in a pool. When a new application requests a connection, DB2 uses one from the existing pool. Using the existing connection reduces the total amount of connection time and the high processor connection cost on the host.

For the data warehouse, connection pooling is implemented using the Tivoli Monitoring environment variable `KHD_CNX_POOL_SIZE`. The default value is 10. You can use the volume of work that the database is processing to decide whether to increase or decrease this value. This parameter is applicable to other database managers such as SQL Server and Oracle.

For information about other environment variables, see "Warehouse Proxy Agent" on page 29 and "Warehouse Summarization and Pruning Agent" on page 30.

Hardware design and operating system usage

For any database system, a number of common areas must be addressed and sized appropriately to support your application workload. This section covers common and platform-specific hardware and operating system components and explains main considerations for hardware design and operating system usage.

It does not include detailed calculations for capacity planning purposes.

Key factors for hardware design and operating system usage include memory, CPU, I/O and network considerations.

Memory:

Understanding how DB2 organizes memory helps you tune memory use for good performance. Many configuration parameters affect memory usage. Some parameters might affect memory on the server, some on the client, and some on the server and the client. Furthermore, memory is allocated and de-allocated at different times and from different areas of the system.

While the database server is running, you can increase or decrease the size of memory areas inside the database shared memory. Understand how memory is divided among the different heaps before you tune to balance overall memory usage on the entire system. See the *DB2 Administration Guide: Performance* for a detailed explanation of the DB2 memory model and all of the parameters that affect memory usage.

CPU:

CPU utilization should be about 70 to 80% of the total CPU time. Lower utilization means that the CPU can cope better with peak workloads. Workloads between 85% to 90% result in queuing delays for CPU resources, which affect response times. CPU utilization above 90% usually causes unacceptable response times.

While running batch jobs, backups, or loading large amounts of data, the CPU might be driven to high percentages such as to 80 to 100% to maximize throughput.

DB2 supports the following processor configurations:

Uni-Processor

A single system that contains only one single CPU

SMP (symmetric multiprocessor)

A single system that can contain multiple CPUs. Scalability is limited to the CPU sockets on the motherboard.

MPP (massively parallel processors)

A system with multiple nodes connected over a high speed link. Each node has its own CPUs. Adding new nodes achieves scalability.

Note:

1. Inefficient data access methods cause high CPU utilization and are major problems for database system. Regular database maintenance is an important factor.
2. Paging and swapping require CPU time. Consider this factor while planning your memory requirements.

I/O: Improving I/O can include making accurate calculations for total disk space that an application requires, improving disk efficiency, and providing for parallel I/O operations.

Calculating disk space for an application:

Use the guidelines in this section to calculate total disk space that an application requires.

- Calculate the raw data size. Add the column lengths of your database tables and multiply by the number of expected rows.
- After calculating the raw data size, use the following scaling up ratios to include space for indexing, working space, and so on:
 - OLTP ratio: 1:3
 - DSS ratio: 1:4
 - Data warehouse ratio: 1:5

Disk efficiency:

You can improve disk efficiency with attention to the concerns outlined here.

- Minimize I/O. Access to main memory is much faster than accessing the disk. Provide as much memory as possible to the database buffer pools and various memory heaps to avoid I/O.
- When I/O is needed, reading simultaneously from several disks is the fastest method. You can use several smaller disks rather than one large disk, or place the disk drives on separate controllers.

Selecting disk drives:

Disks tend to grow larger every year, doubling in capacity every 18 months, and the cost per GB is lower each year. The cost difference of the two smallest drives diminishes until the smaller drive is not practical. The disk drives improve a little each year in seek time. The disk drives get smaller in physical size. While the disk drives continue to increase capacity with a smaller physical size, the speed improvements, seek, and so on, are small in comparison. A database that would have taken 36 * 1 GB drives a number of years ago can now be placed on one disk.

This growth trend highlights database I/O problems. For example, if each 1 GB disk drive can do 80 I/O operations a second, the system can process a combined 2880 I/O operations per second (36 multiplied by 80). But a single 36-GB drive with a seek time of 7 milliseconds can process only 140 I/O operations per second. Although increasing disk drive capacity is beneficial, fewer disks cannot deliver the same I/O throughput.

Parallel operations:

Provide for parallel I/O operations. Use the smallest disk drives possible to increase the number of disks for I/O throughput.

If you purchase larger drives, use only half the space, especially the middle area, for the database. The other half is useful for backups, archiving data, off hour test databases, and extra space for accommodating upgrades.

Network:

After a system is implemented, the network should be monitored to ensure that its bandwidth is not being consumed more than 50%.

The network can influence overall performance of your application, especially if a delay occurs in the following situations:

- Lengthy time between the point a client machine sends a request to the server and the server receives this request
- Lengthy time between the point the server machine sends data back to the client machine and the client machine receives the data

Tuning

This section explains relevant database and database manager configuration parameters and includes guidelines for setting their values.

Review these topics and considerations for relevancy to your environment.

Database manager configuration tuning:

Each instance of the database manager has a set of database manager configuration parameters, also called *database manager parameters*, which affect the amount of system resources that are allocated to a single instance of the database manager.

Some of these parameters are used for configuring the setup of the database manager and other information that is not related to performance. You can use either the DB2 Control Center or the following command to change the parameters:

```
UPDATE DATABASE MANAGER CONFIG USING keyword
```

Parameters:

The database manager configuration parameters in this section have a high impact on performance.

Note: See the *DB2 Administration Guide: Performance* for detailed explanations about all the database manager configuration parameters.

AGENTPRI

Controls the priority that the operating system scheduler gives to all agents and to other database manager instance processes and threads. Use the default value unless you run a benchmark to determine the optimal value.

ASLHEAPSZ

Represents a communication buffer between the local application and its associated agent. The application support layer heap buffer is allocated as shared memory by each database manager agent that is started. If the request to the database manager or its associated reply do not fit into the buffer, they are split into two or more send-and-receive pairs. The size of this buffer should be set to handle the majority of requests using a single send-and-receive pair.

INTRA_PARALLEL

Specifies whether the database manager can use intra-partition parallelism on an SMP machine. Multiple processors can be used to scan and sort data for index creation. Usually, this parameter is set to YES, especially if you are running on a dedicated database server. The default value is NO.

MAX_QUERYDEGREE

Specifies the maximum degree of intra-partition parallelism that is used for any SQL statement that is executing on this instance of the database manager. An SQL statement does not use more than this number of parallel operations within a partition when the statement is executed. The `intra_parallel` configuration parameter must be set to a value greater than 1 to enable the database partition to use intra-partition parallelism. Setting the value to ANY enables use of all partitions. Usually, this parameter should be set to ANY, especially if you are running on a dedicated database server. The default value is ANY.

SHEAPTHRES

Determines the maximum amount of memory available for all the operations that use the sort heap, including sorts, hash joins, dynamic bitmaps that are used for index ANDing and Star Joins, and operations where the table is in memory. Set the sort heap threshold parameter to a reasonable multiple of the largest SORTHEAP parameter in your database manager instance. This parameter should be at least twice as large as the largest sort heap defined for any database within the instance, but you must also consider the number of concurrent sort processes that can run against your database.

Database configuration tuning:

Each database has a set of the database configuration parameters, which are also known as *database parameters*. These parameters affect the amount of system resources that are allocated to that database. Furthermore, some database configuration parameters provide descriptive information only and cannot be changed, and others are flags that indicate the status of the database.

You can use the DB2 Control Center or the `UPDATE DATABASE CONFIG FOR dbname USING keyword` command to change those parameters. For more information about the numerous database configuration parameters, see the *DB2 Administration Guide: Performance*.

The following data configuration parameters have a high impact on performance:

DBHEAP

Contains control block information for tables, indexes, table spaces, and buffer pools, and space for the log buffer (LOGBUFSZ) and temporary memory that the utilities use. Each database has only one database heap, and the database manager uses it on behalf of all applications connected to the database. The size of the heap is dependent on a large number of variables. The control block information is kept in the heap until all applications disconnect from the database. The DB2 default value is typically too low, particularly for the Tivoli Data Warehouse. Start with a value between 2000 and 8000.

DFT_DEGREE

Specifies the default value for the CURRENT DEGREE special register and the DEGREE bind option. The default value is 1, which means no intra-partition parallelism. A value of -1 means that the optimizer determines the degree of intra-partition parallelism based on the number of processors and the type of query. The degree of intra-partition parallelism for an SQL statement is specified at statement compilation time using the CURRENT DEGREE special register or the DEGREE bind option. The maximum runtime degree of intra-partition parallelism for an active application is specified using the **SET RUNTIME DEGREE** command. The Maximum Query Degree of Parallelism (max_querydegree) configuration parameter specifies the maximum query degree of intra-partition parallelism for all SQL queries. The actual runtime degree that is used is the lowest of one of the following:

- The max_querydegree configuration parameter
- Application runtime degree
- SQL statement compilation degree

For a multi-processor machine, set this to -1 (ANY), to allow intra-partition parallelism for this database.

CHNGPGS_THRESH

Improves overall performance of the database applications. Asynchronous page cleaners write changed pages from the buffer pool or the buffer pools to disk before a database agent requires the space in the buffer pool. As a result, database agents should not have to wait for changed pages to be written out so that they might use the space in the buffer pool. Usually, you can start out with the default value.

LOCKLIST

Indicates the amount of storage that is allocated to the lock list. This parameter has a high impact on performance if frequent lock escalations occur. The DB2 default value is typically too low, particularly for the Tivoli Data Warehouse. Start with a value of between 500 and 800.

MAXLOCKS

Maximum percent of lock list before escalation. Use this parameter with the LOCKLIST parameter to control lock escalations. Increasing the LOCKLIST parameter augments the number of available locks.

LOGBUFSZ

Specifies the amount of the database heap (defined by the dbheap parameter) to use as a buffer for log records before writing these records to disk. Buffering the log records supports more efficient logging file I/O because the log records are written to disk less frequently, and more log records are written at each time. The DB2 default value is typically too low, particularly for the Tivoli Data Warehouse. Start with a value of between 256 and 768.

NUM_IOCLEANERS

Specifies the number of asynchronous page cleaners for a database. These page cleaners write changed pages from the buffer pool to disk before a database agent requires the space in the buffer pool. As a result, database agents do not wait for changed pages to be written out so that they can use the space in the buffer pool. This parameter improves overall performance of the

database applications. The DB2 default value is typically too low. Set this parameter equal to the number of physical disk drive devices that you have. The default value is 1.

NUM_IOSERVERS

Specifies the number of I/O servers for a database. I/O servers perform prefetch I/O and asynchronous I/O by utilities such as backup and restore on behalf of the database agents. Specify a value that is one or two more than the number of physical devices on which the database is located. The DB2 default value is typically too low. Set this parameter equal to the number of physical disk drive devices that you have, and add two to that number. The default value is 3.

PCKCACHESZ

The package cache is used for caching sections for static and dynamic SQL statements on a database. Caching packages and statements eliminates the requirement to access system catalogs when reloading a package so that the database manager can reduce its internal overhead. If you are using dynamic SQL, caching removes the need for compilation.

SORTHEAP

Defines the maximum number of private memory pages to be used for private sorts, or the maximum number of shared memory pages to be used for shared sorts. Each sort has a separate sort heap that is allocated as needed by the database manager. This sort heap is the area where data is sorted. Increase the size of this parameter when frequent large sorts are required. The DB2 default value might be too low, particularly for the Tivoli Data Warehouse. Start with a value of 256 - 1024. When changing this parameter, you might want to change the SHEAPTHRES database manager parameter too.

LOGFILSIZ

Defines the size of each primary and secondary log file. The size of these log files limits the number of log records that can be written to them before they become full, which requires a new log file. The DB2 default value is too low, particularly for the Tivoli Data Warehouse. Start with a value of 4096 - 8192.

LOGPRIMARY

Specifies the number of primary log files to be pre-allocated. The primary log files establish a fixed amount of storage that is allocated to the recovery log files. The DB2 default value might be too low, particularly for the Tivoli Data Warehouse. Start with a value of 6 - 10.

Buffer pools:

The buffer pool is the area of memory where database pages, table rows or indexes, are temporarily read and manipulated. All buffer pools are located in global memory, which is available to all applications using the database. The purpose of the buffer pool is to improve database performance.

Data can be accessed much faster from memory than from disk. Therefore, as the database manager is able to read from or write to more data (rows and indexes) to memory, database performance improves.

The default buffer pool allocation is usually not sufficient for production applications, and must be monitored and tuned before placing your application in production. The DB2 default value is typically too low, particularly for the Tivoli Data Warehouse. Start with a value of 50 -75 percent of your system memory if the database server is dedicated. You can use the SQL statement ALTER BUFFERPOOL to change this value.

Registry variables:

Each instance of the database manager has a set of registry and environment variables that affect various aspects of DB2 processing. You can change the value of DB2 registry variables using the **DB2SET** command. Although numerous other registry and environment variables exist, the DB2_PARALLEL_IO variable has a high impact on performance.

Note: See the *DB2 Administration Guide: Performance* for a detailed explanation of all the registry and environment variables.

While reading or writing data from and to table space containers, DB2 may use parallel I/O for each table space value that you specify. The prefetch size and extent size for the containers in the table space determine the degree of parallelism. For example, if the prefetch size is four times the extent size, there are four extent-sized prefetch requests. The number of containers in the table space does not affect the number of prefetchers.

To enable parallel I/O for all table spaces, you can specify the asterisk (*) wildcard character. To enable parallel I/O for a subset of all table spaces, enter the list of table spaces. For several containers, extent-size pieces of any full prefetch request are broken down into smaller requests that are executed in parallel based on the number of prefetchers. When this variable is not enabled, the number of containers in the table space determines the number of prefetcher requests that are created.

Monitoring tools

This section lists the tools that DB2 provides for monitoring or analyzing your database.

- Snapshot Monitor, which captures performance information at periodic points of time and is used to determine the current state of the database
- Event Monitor, which provides a summary of activity at the completion of events such as statement execution, transaction completion, or when an application disconnects
- Explain Facility, which provides information about how DB2 accesses the data to resolve the SQL statements
- db2batch tool, which provides performance information as a benchmarking tool

SNAPSHOT and EVENT monitors:

DB2 maintains data as the database manager runs about its operation, its performance, and the applications that are using it. This data can provide important performance and troubleshooting information.

For example, you can track the following developments:

- The number of applications connected to a database, their status, and which SQL statements each application is executing, if any
- Information that shows how well the database manager and database are configured, helping you make tuning decisions
- Information about the time that deadlocks occurred for a specified database, the applications that were involved, and the locks that were in contention
- The list of locks held by an application or a database. If the application cannot proceed because it is waiting for a lock, additional information is on the lock, including which application is holding it.

Collecting performance data introduces some overhead on the operation of the database. DB2 provides monitor switches to control which information is collected. You can use the following DB2 commands to turn on these switches:

```
UPDATE MONITOR SWITCHES USING BUFFERPOOL ON ;
UPDATE MONITOR SWITCHES USING LOCK ON ;
UPDATE MONITOR SWITCHES USING SORT ON ;
UPDATE MONITOR SWITCHES USING STATEMENT ON ;
UPDATE MONITOR SWITCHES USING TABLE ON ;
UPDATE MONITOR SWITCHES USING UOW ON ;
```

You can access the data that the database manager maintains either by taking a snapshot or by using an event monitor.

SNAPSHOTs:

Use the **GET SNAPSHOT** command to collect status information and format the output for your use. The returned information represents a snapshot of the database manager operational status at the time the command was issued.

Various formats of this command obtain different kinds of information, and the specific syntax can be obtained from the DB2 Command Reference. The following formats are useful:

GET SNAPSHOT FOR DATABASE

Provides general statistics for one or more active databases on the current database partition.

GET SNAPSHOT FOR APPLICATIONS

Provides information about one or more active applications that are connected to a database on the current database partition.

GET SNAPSHOT FOR DATABASE MANAGER

Provides statistics for the active database manager instance.

GET SNAPSHOT FOR LOCKS

Provides information about every lock held by one or more applications connected to a specified database.

GET SNAPSHOT FOR BUFFERPOOLS

Provides information about buffer pool activity for the specified database.

GET SNAPSHOT FOR DYNAMIC SQL

Returns a point-in-time picture of the contents of the SQL statement cache for the database.

You can create some simple scripts and schedule them to get periodic snapshots during your test cycles.

DB2BATCH:

A benchmark tool called DB2BATCH is provided in the sqllib/bin subdirectory of your DB2 installation. This tool can read SQL statements from either a flat file or standard input, dynamically describe and prepare the statements, and return an answer set.

You can specify the level of performance information that is supplied, including the elapsed time, CPU and buffer pool usage, locking, and other statistics collected from the database monitor. If you are timing a set of SQL statements, DB2BATCH also summarizes the performance results and provides both arithmetic and geometric means. For syntax and options, type db2batch.

Optimizing queries

This section contains information about tuning the queries that are processed to display the tables, charts, and graphs that make up workspace views within the Tivoli Enterprise Portal.

Review these topics and considerations for relevancy to your environment.

Processing queries

The query assigned to a chart or table view requests data from a particular attribute group. It executes when you open or refresh the workspace. Queries make up the processing load for on-demand data collection.

You can reduce the frequency and amount of data sampling by:

- Customizing the query to filter out unwanted data. This reduces the number of selection criteria (rows) and attributes (columns) collected.

- Applying the same query to other views in the workspace. This reduces the number of data samples required: one query uses a single sample for multiple views.
- Disabling automatic refresh of workspace views or adjust the refresh rate to longer intervals. This causes Tivoli Enterprise Monitoring Agent data to be collected less frequently.
- Consider how you want to display the data returned by a query. A graphical view workspace might require less data compared to a table view, because it uses only the truly numeric data and leaves out the character data.

Do not confuse custom queries with view filters from the **Filters** tab of the Query Properties editor. View filters fine-tune the data after it has been retrieved by the query and do not reduce network traffic, data collection processing, or memory demands.

The following general recommendations and observations might be considered as well:

- Some attributes are more expensive to retrieve than others. One *expensive* column in a table makes any workspace view or situation that references that table more expensive. An example of an expensive attribute is one that must run long storage chains to determine its value, such as using a process table to look for looping tasks. Where possible, ensure that you only retrieve attributes that are required for the monitoring process.
- Column function (such as MIN, MAX, AVG, and so on) requires post processing of the query results set after data is returned to Tivoli Enterprise Monitoring Server.
- Use more efficient data manipulating functions, such as substring instead of string scan. If you know the position of the string to search, do not scan the whole string to check for the value.

Post-filters versus pre-filters

Performance is improved for each view if you pre-filter the required view information and send to the portal client only what is needed in that view. This reduces the amount of data sent through the network and you do not have to post-process it either. However, there is one exception to the rule.

Think of a workspace that contains many views. Each of those views has a query associated with it, which is issued when the workspace is accessed. This might result in many queries that must be processed in parallel.

A better way of doing this might be to create one query that returns all data that is required in the workspace. In this case, the query will only be issued once and the data can then be post-filtered for each view to only display information as it applies to each view.

One important consideration however is that queries are saved globally and are not user ID dependent. This means that only administrators will be able to modify queries in most installations. For the end user to be able to modify filters, the preferred method might therefore be the filters applied in the view properties **Filters** tab.

Defining custom queries

Custom queries reduce network traffic, processing at the agent and Tivoli Enterprise Monitoring Server, and memory usage at the Tivoli Enterprise Portal Server and Tivoli Enterprise Portal client. Custom queries accomplish this by limiting the number of rows and columns passed from the Tivoli Enterprise Monitoring Agent to the Tivoli Enterprise Monitoring Server.

Most of the predefined, predefined queries request all columns and all rows of an attribute group, of which only a few might be of interest to you. Removing the unwanted columns (attributes) reduces the amount of data transferred from monitoring agent to Tivoli Enterprise Monitoring client via the portal server, hub monitoring server, and remote monitoring server. Additionally, in the case of OMEGAMON Monitoring Agents that are located on z/OS, you also reduce the amount of EBCDIC/ASCII conversion required when data is passed between mainframe and distributed platforms.

It is recommended to tune any queries servicing workspaces that are frequently executed or return large quantities of data. Query execution always requires resources and intermittent large reports causes a spike in memory requirements and network resources.

Restricting the number of rows

Most predefined queries return all rows and columns. You can create a custom query to filter out the irrelevant or uninteresting metrics. Not only does this make it easier to read the report, but it saves Tivoli Enterprise Portal Server memory, client memory, and CPU because there are fewer rows to translate, sort, and transmit.

You might be using view filters inappropriately. View filters work only on the current page returned from the query. For example, if page size is 100 lines and the filter reduces it to five lines on page 1 and similarly on subsequent pages, the row set cannot be seen on one page. Do not increase the workspace page size to see everything on one page. Increased page size actually increases portal client memory requirements.

Instead, avoid this condition by creating a custom query that filters the data at query execution time.

Restricting the number of columns

Most predefined queries return all columns (or attributes). A particular attribute group might contain 50 columns, yet all you need is five. Creating a custom query to retrieve only the required five attributes reduces portal server and client CPU and memory.

Use the same query in a workspace

If you have multiple views in a workspace requiring data from different attribute groups, you will need a different query for each group. However, if the views have data from the same attribute group, then use a single query to accommodate both, even if the attributes (columns) that each view requires are different. Two unique queries will each drive data collection at the agent and increase resource overhead. For each workspace, have one query that can be shared by all views using that attribute group. Remember that the entire results set for each query is stored on the portal server, this technique avoids duplicate result sets being stored.

Collect agent data less frequently

A good practice is to avoid the use of workspace automatic refresh when possible. The Navigator view and event views (message log, event console, and graphic console) refreshes automatically. This provides you with instantaneous alerts, and you can navigate to their event workspaces with actual data. The graphic view, the graphical equivalent of the Navigator, which shows alerts but no data, affects client memory but not the portal server.

Reduce the distribution of queries

A query can be assigned to a particular managed system or to a group of systems by defining a managed system group. The default group usually includes all known instances for that application or operating system.

It might be that you want this query to be applied to only certain managed systems and so distribution to all managed systems is an unnecessary waste of resources. Modify MSLs to reduce the distribution of the query. Also remove the MSLs for queries that are of no interest to the user. Even if you are not viewing the results of the query, there might be a use of system resources that can be avoided by restricting the distribution of the unneeded queries.

Optimizing situations

Situations are conditions that you want to monitor on managed systems. Each situation contains a name, a predicate formula, special attributes for specific situation processing, information about whether the situation is automatically started or not, and the sampling interval. It can also contain a command to execute when the situation is true, and advice to give the client when an alert for the situation is surfaced, and so on.

A situation predicate is a formula that uses attributes, functions, and other situations as operands along with threshold values to describe a filtering criterion. The situation formula is made up of one or more logical expressions.

Each expression takes the form:

[Attribute name] / [logical operator] / [value]

For example: `PROCESS_ID == 0`

The situation predicates are similar to a WHERE clause in SQL. In IBM Tivoli Monitoring, predicates are processed sequentially, and perform better if you put the most restrictive condition as the first predicate. You can use the situation editor to create/edit situations by choosing attributes, logical operator, value, and sampling interval, and so on. Situations are assigned to run on behalf of managed systems or lists of managed systems.

When a situation is running on a monitoring agent, the agent collects the current values of the attributes specified in the formula, and tests the values against a threshold. When the condition is met, that is the threshold is exceeded or a value is matched, the agent passes the collected data back to the monitoring server to which it is connected, and an event is generated.

However, some types of situations cannot be tested at the agent level:

- Situations involving group functions, such as MIN, MAX, AVG, SUM, COUNT
- Embedded situations
- Correlated situations

For these types of situations, the agent returns all rows back to the monitoring server to which it is connected, and the server performs the testing. In large-scale environments, especially if the sampling interval for the situation is short, evaluation for such situations dramatically increases the workload and memory usage for the monitoring server.

In general, there is limited situation processing at the hub monitoring server if there are no agents directly connected to the hub. For remote monitoring servers, there is a direct correlation to the number of agents, situations, and data rows to the amount of memory required. Therefore, the number of situations and size of the data might be the limiting factor that determines how many agents an individual remote monitoring server can support.

Since the performance of the monitoring server is greatly affected by the volume and frequency of situation state changes, do not run a situation and collect data unless you are interested in the results.

The following recommendations provide guidance for how to write more efficient situations and how to reduce the situation processing requirements:

1. If possible, avoid use of group functions, embedded situations, or correlated situations. Processing demands for such situations are much higher, since all attribute group rows must be sent to the monitoring server for testing, increasing the processing demand and memory usage on the monitoring server.
2. Put the most stringent condition at the beginning of the formula because the conditions are evaluated sequentially, left to right.

Consider a situation that has the first condition test on real storage use, the result set may contain multiple rows; then the second condition tests whether a given process name is among the returned rows. It would be more efficient to first test on process name (the result will be one row), followed by the test on the storage usage, just on the single row.

Consider the following in creating the condition tests:

- a. Numeric attributes are processed more quickly than text attributes.

- b. String checking with substring (STR) is more efficient than the string scan (SCAN), especially for long strings. If you know the exact location of the text or characters to be evaluated, use a substring.
3. Use longer sampling intervals where possible, especially for situations that are distributed to many agents.
4. Minimize the number of situations on attribute groups that can return many rows, such as process or disk attribute groups.
5. Put mission critical systems in a separated managed system group, and distribute the heavy workload situations to them only if necessary.
6. Consider spreading business-critical situations among several different remote Tivoli Enterprise Monitoring Servers. For example, a database monitoring situation might be high load, but business-critical. Instead of having all 500 database agents report through the same remote monitoring server, consider configuring the database agents across five remote monitoring servers, which are supporting other less-demanding agents.

Validating your installation

You can check the release version of the common installer used for your currently installed image of Tivoli Monitoring. The installer version is logged as the first text line of the main installation or upgrade log file and it should read, for example, *Installation Logging for IBM Tivoli Monitoring Version 6.3.0*.

You can also validate whether specific binary KinConfig.exe, ITMISUtil.dll, and ServiceStatus.exe files have Tivoli Monitoring Version 6.3.0 by right-clicking files in Windows Explorer and reviewing the version properties tab.

You can use the **kininfo** command to validate your installation. On Windows systems, view information for your monitoring server, including inventory of installed IBM Tivoli products, configuration settings, installed CD versions, and a list of running IBM Tivoli processes.

CLI syntax

```
kininfo
[-d]
[-i]
[-r]
[-l]
[-t]
```

Where:

- **-d** Displays a list of installed products, which can be parsed.
- **-i** Lists the inventory in English.
- **-r** Displays a list of running agents.
- **-l** Displays the log switch.
- **-t** Displays the product name, version, build information, and installation date for all of the products installed in the installation directory.

You can also use this option to review the installed support of self-described agents, which is displayed in a table. The following example shows the table for the r2 agent:

```
***** Thursday, May 20, 2010 3:23:29 PM *****
User : Administrator Group : NA
Host Name : NC045161 Installer : Ver: 06300000
CandleHome : C:\IBM\ITM
Installitm : C:\IBM\ITM\InstallITM
*****
...Product Inventory
PC PRODUCT DESC PLAT VER
BUILD INSTALL DATE
```

```

R2 Agentless Monitoring for Windows Operating Sy WINNT 06.30.00.00
201004121113 20100520 1459
PC SELF-DESCRIBED APPLICATION SUPPORT PACKAGE PLAT APP VER
R2 Agentless Monitoring for Wind - r2tms_support CMS 06.30.00.00
R2 Agentless Monitoring for Wind - r2tps_support CNS 06.30.00.00
R2 Agentless Monitoring for Wind - r2tpw_support XEB 06.30.00.00
PC APPLICATION SUPPORT DESC PLAT APP VER
BUILD INSTALL DATE
R2 Agentless Monitoring for Windows Operating Sy WICMS 06.30.00.00
201004121113 20100520 1512

```

CLI example

The following example shows all installed products.

```
kincinfo -i
```

The following is the output of this example:

```

kincinfo -i ***** Mon May 07 14:19:20 Eastern Daylight Time 2007 *****
User : Administrator Group : NA
Host Name : FVWIN18 Installer: Ver: 0NOVALUE00000
CandleHome: C:\IBM\ITM
*****
...Product Inventory
A4 i5/OS Support
WINNT Version: 06.10.05.01 Build: 200702230014
A4 i5/OS Support
WINNT Version: 06.10.05.01 Build: 200702230014
A4 i5/OS Support
WINNT Version: 06.10.05.01 Build: 200702230014
A4 i5/OS Support
WINNT Version: 06.10.05.01 Build: 200702230014
AX Tivoli Enterprise Monitoring Agent Framework
WINNT Version: 03.50.03.00 Build: 200510061051
CJ Tivoli Enterprise Portal Desktop Client
WINNT Version: 06.10.05.01 Build: 200705012123
CQ Tivoli Enterprise Portal Server
WINNT Version: 06.10.05.01 Build: 200705012135
CW Tivoli Enterprise Portal Browser Client
WINNT Version: 06.10.05.01 Build: 200705012123
GL Tivoli Enterprise Monitoring Agent Framework
WINNT Version: 06.10.05.01 Build: 200705012139
HD Warehouse Proxy
WINNT Version: 06.10.05.01 Build: 200705012139
IT TEC GUI Integration
WINNT Version: 06.10.05.01 Build: 200611010030
IT TEC GUI Integration
WINNT Version: 06.10.05.01 Build: 200611010031
IT TEC GUI Integration
WINNT Version: 06.10.05.01 Build: 200611010030
KF IBM Eclipse Help Server
WINNT Version: 06.10.05.01 Build: 200704171107
LZ Linux OS Support
WINNT Version: 06.10.05.01 Build: 200704301644
LZ Linux OS Support
WINNT Version: 06.10.05.01 Build: 200704301644
LZ Linux OS Support
WINNT Version: 06.10.05.01 Build: 200704301644
LZ Linux OS Support
WINNT Version: 06.10.05.01 Build: 200704301644
MS Tivoli Enterprise Monitoring Server
WINNT Version: 06.10.05.01 Build: 200705012139
NS NLS Support
WINNT Version: 03.50.03.00 Build: 200510061051
NT Monitoring Agent for Windows OS
WINNT Version: 06.10.05.01 Build: 200701160818
NT Windows OS Support

```

WINNT Version: 06.10.05.01 Build: 200701160818
NT Windows OS Support
WINNT Version: 06.10.05.01 Build: 200701160818
NT Windows OS Support
WINNT Version: 06.10.05.01 Build: 200701160818
NT Windows OS Support
WINNT Version: 06.10.05.01 Build: 200701160818
SY Summarization and Pruning agent
WINNT Version: 06.10.05.01 Build: 200705012135
SY Summarization and Pruning agent
WINNT Version: 06.10.05.01 Build: 200705012135
TM IBM Tivoli Monitoring 5.x Endpoint Support
WINNT Version: 06.10.05.01 Build: 200604051327
TM IBM Tivoli Monitoring 5.x Endpoint Support
WINNT Version: 06.10.05.01 Build: 200604051327
UI Tivoli Enterprise Services User Interface
WINNT Version: 06.10.05.01 Build: 200705012139
UL UNIX Logs Support
WINNT Version: 06.10.05.01 Build: 200701120847
UL UNIX Logs Support
WINNT Version: 06.10.05.01 Build: 200701120847
UL UNIX Logs Support
WINNT Version: 06.10.05.01 Build: 200701120847
UL UNIX Logs Support
WINNT Version: 06.10.05.01 Build: 200701120847
UM Universal Agent
WINNT Version: 06.10.05.01 Build: 200705012131
UM Universal Agent Support
WINNT Version: 06.10.05.01 Build: 200705012131
UM Universal Agent Support
WINNT Version: 06.10.05.01 Build: 200705012131
UM Universal Agent Support
WINNT Version: 06.10.05.01 Build: 200705012131
UM Universal Agent Support
WINNT Version: 06.10.05.01 Build: 200705012131
UX UNIX OS Support
WINNT Version: 06.10.05.01 Build: 200701161329
UX UNIX OS Support
WINNT Version: 06.10.05.01 Build: 200701161329
UX UNIX OS Support
WINNT Version: 06.10.05.01 Build: 200701161329
UX UNIX OS Support
WINNT Version: 06.10.05.01 Build: 200701161329
C:\ >

Chapter 19. Tivoli Data Warehouse solutions

The term *Tivoli Data Warehouse solution* refers to the set of IBM Tivoli Monitoring components, successfully installed and configured, that interact to collect and manage historical data.

These *warehousing components* include the Tivoli Enterprise Monitoring Server and the Tivoli Enterprise Portal Server, the Tivoli Data Warehouse database, the Warehouse Proxy Agent, the Summarization and Pruning Agent, and the Schema Publication Tool.

While all Tivoli Data Warehouse solutions include these six components, the details of any particular solution vary according to the size of the environment, which database management systems are used, and which operating systems are involved. This chapter introduces and summarizes your options for database platforms, operating systems, and communications between warehousing components.

- “Planning assumptions” on page 515
- “Firewall considerations for the Tivoli Data Warehouse” on page 517
- “Summary of supported operating systems” on page 519

Each of the next four chapters describes how to implement solutions for a variety of configurations using one of the supported database platforms: IBM DB2 for Linux, UNIX, and Windows, IBM DB2 on z/OS, Microsoft SQL Server, and Oracle.

New in Version 6.3

See the following changes to the Tivoli Data Warehouse component for Tivoli Monitoring V6.3.

Tivoli Data Warehouse range partitioning

Range partitioning is a database data organization feature that can significantly improve pruning and query performance in large Tivoli Data Warehouse databases. For more information and the associated configuration parameters, see Chapter 21, “Tivoli Data Warehouse range partitioning,” on page 531.

Specify the tablespaces where tables and indices are created

Beginning with Version 6.3, you can configure the Warehouse Proxy Agent and the Summarization and Pruning Agent so that created Tivoli Data Warehouse tables and indices are located in default table spaces. Tablespaces can be specified for partitioned or non-partitioned tables. This new function is independent of partitioning. For more information, see Chapter 22, “Specifying where tables and indices are created,” on page 537.

New Summarization and Pruning Agent workspace

A new workspace, Table Partition Statistics, is available from the Statistics Navigator item. This workspace provides statistics about the table partitions if table range partitioning is enabled. You can see the number of partitions created, added, removed, or rotated per table per product.

Migrating non-partitioned tables to partitioned tables using the Schema Publication Tool

The schema publication tool allows you to generate SQL scripts with tables and indices created in default table spaces. This new implementation in Tivoli Monitoring V6.3 provides a means for you to migrate existing non-partitioned tables to partitioned tables. This feature allows you to benefit from the performance improvements of partitioned tables. For more information, see *IBM Tivoli Monitoring*

Administrator's Guide (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/adminuse/itm_admin.htm).

Filter by table in Schema Publication Tool

You can now filter by table in any mode of the Schema Publication Tool. For more information, see “Generating SQL for data warehouse tables” on page 525.

Planning considerations for the Tivoli Data Warehouse

In a large or enterprise environment, you have the option of using multiple databases for the Tivoli Data Warehouse or having all of your hub monitoring servers use one single database. The benefit of using one database across your environment is that the information is stored in one location, meaning that you can more easily and accurately generate reports that reflect your entire environment (as opposed to needing to collate reports from several different databases).

However, because the amount of data that is generated during history collection across a large or enterprise environment can be immense, you need to carefully plan the scale and performance of that database.

In addition to the planning information below, the IBM redbook *Tivoli Management Services Warehouse and Reporting*, SG24-7290, provides database performance tuning information. You can download the book at the following location: <http://www.redbooks.ibm.com/abstracts/sg247290.html?Open>.

Estimating the required size of your database

One of the factors to consider when planning the size of database that you need is the amount and type of information you will collect for agent history data collection. Each of the agent user's guides provide capacity planning information to help you calculate the amount of disk space required by data for each attribute group.

Use the information in this section to complete the calculations for determining how large your data warehouse database needs to be.

- “Step 1: Determine the number of detailed records per day for each attribute group”
- “Step 2: Determine the hard disk drive footprint for each attribute group” on page 511
- “Step 3: Determine the amount of detailed data for each attribute group” on page 511
- “Step 4: Calculate the amount of aggregate data for each attribute group” on page 511
- “Step 5: Determine the estimated size of your database” on page 512

Use the worksheets in Appendix A, “Installation worksheets,” on page 833 to record the values from these calculations. To print these, go to the IBM Tivoli Monitoring information center: IBM Tivoli Monitoring Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/welcome.htm).

You might use the Warehouse Load Projection tool available in the IBM Tivoli Integrated Service Management Library. (Search on “Warehouse Load Projects at the following site: IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>)).) The tool does all the calculations for you and includes data for nearly all of the IBM Tivoli Monitoring V6.x-based monitoring agents.

Step 1: Determine the number of detailed records per day for each attribute group

Determine the number of detailed records per day for each attribute group that you want to collect data for.

Use the following equation:

$(60 / \text{collection interval}) * (24) * (\# \text{ instances at each interval})$

where:

60 Represents the 60 minutes in an hour.

collection interval

The data collection interval, in minutes. This value can be 1, 5, 15, 30, 60, or 1440 (1 day).

24 Represents 24 hours in one day.

instances at each interval

The number of instances recorded at each interval. See the agent user's guide for this value.

Step 2: Determine the hard disk drive footprint for each attribute group

Determine the hard disk drive footprint for each attribute group. The result generated by this formula gives an estimate of the amount of disk space used for this attribute group for 24 hours worth of data for a single agent.

Use the following equation:

$(\# \text{ detailed records}) * (\text{attribute group detailed record size}) / 1024$

where:

detailed records

The number of detailed records for the attribute. This is the value you calculated in "Step 1: Determine the number of detailed records per day for each attribute group" on page 510.

attribute group detailed record size

The detailed record size for the attribute group. See the agent user's guide for this value.

1024 Represents 1 KB and causes the equation to generate a kilobyte number instead of a byte number.

Step 3: Determine the amount of detailed data for each attribute group

Determine the amount of detailed data in the warehouse database for each attribute group.

Use the following equation:

$(\text{attribute group disk footprint}) * (\# \text{ of agents}) * (\# \text{ days of detailed data}) / 1024$

where:

attribute group disk footprint

The disk footprint for the attribute group. This is the value you calculated in "Step 2: Determine the hard disk drive footprint for each attribute group."

of agents

The number of agents of the same agent type in your environment.

days of detailed data

The number of days for which you want to keep detailed data in the warehouse database.

1024 Represents 1 KB and causes the equation to generate a megabyte number.

Step 4: Calculate the amount of aggregate data for each attribute group

Determine the amount of aggregate data in the warehouse for each attribute group.

First, calculate the number of aggregate records per agent. Use the following equation:

$(\# \text{ hourly} + \# \text{ daily} + \# \text{ weekly} + \# \text{ monthly} + \# \text{ quarterly} + \# \text{ yearly}) * (\# \text{ instances at each interval})$

where:

#hourly

The number of hourly records for the attribute group. For example, if you have hourly records for 60 days, the number of hourly records is 1440 (60 multiplied by 24 hours per day).

#daily The number of daily records for the attribute group. For example, if you have daily records for 12 months, the number of daily records is 365.

#weekly

The number of weekly records for the attribute group. For example, if you have weekly records for a 2-year period, the number of weekly records is 104 (2 multiplied by 52 weeks per year).

#monthly

The number of monthly records for the attribute group. For example, if you have monthly records for a 2-year period, the number of monthly records is 24.

#quarterly

The number of quarterly records for the attribute group. For example, if you have quarterly records for a 2-year period, the number of quarterly records is 8 (2 years multiplied by 4 quarters in a year).

#yearly The number of yearly records for the attribute group. For example, if you have yearly reports for a 10-year period, the number of yearly records is 10.

instances at each interval

The number of instances recorded at each interval. See the agent user's guide for this value.

Next, use the following equation to calculate the amount of attribute data in the warehouse for the attribute group:

$(\# \text{ aggregate records per agent}) * (\text{attribute group aggregate record size}) * (\# \text{ agents}) / 1048576$

where:

aggregate records per agent

The number of aggregate records per agent for the attribute group.

attribute group aggregate record size

The size of the attribute group aggregate group. See the agent user's guide for this value.

agents

The number of agents of the same agent type in your environment.

1048576

Represents 1 MB and causes the equation to generate a megabyte number.

Step 5: Determine the estimated size of your database

First, determine the total space required for each attribute group. Add the amount of detailed data and the amount of aggregate data for the attribute group.

Use the following equation:

$(\text{detailed data size}) + (\text{aggregate data size})$

Second, determine the total space required for all attribute groups for the agent. Add the total space for each attribute group that you want to collect. Use the following equation:

$\text{aggGroup1} + \text{aggGroup2} + \text{aggGroup3} \dots$

Third, determine the total space required for all agents. Add the total space for each agent. Use the following equation:

$\text{agent1} + \text{agent2} + \text{agent3} \dots$

Use the following worksheet to estimate the size of your data warehouse database.

[illegible]

Consider the factors in this section when designing a disk subsystem to support your database processing.

- The following table provides some example sizes for a database:

Table 80. Database size examples

	Number of disks to use			
	Absolute minimum disks	Small RDBMS	Small and safe RDBMS	Large RDBMS
Operating System	1	1	1 + mirror	1
Paging and RDBMS code	Use above	1	1 + mirror	1
RDBMS data	1	1	1 + mirror	8
RDBMS indexes	1	1	1 + mirror	6
RDBMS temp	Use above	1	1 + mirror	6
RDBMS logs	1 + mirror	1 + mirror	1 + mirror	2
Database data	12 GB	24 GB	48 GB	108+ GB
Number of disks	5	7	12	24

The **Absolute minimum disks** column specifies the minimum number of disks for an RDBMS. In this column, the index and temporary space is allocated onto one disk. While not an ideal arrangement, this might work in practice because databases tend to use indexes for transactions or temporary space for index creation and sorting full table scan large queries, but not both at the same time. This is not a recommended minimum disk subsystem for a database, but it does have the lowest cost.

The **Small RDBMS** column represents a minimum disk subsystem, although there might be limits in I/O rates because of the data being placed on only one disk. Striping the data, indexes, and temporary space across these three disks might help reduce these I/O rate limits. This disk subsystem arrangement does not include disk protection for the database or other disks (apart from the mandatory log disk protection for transaction recovery).

The **Small and safe RDBMS** column adds full disk protection and can withstand any disk crash with zero database downtime.

The **Large RDBMS** column represents a typical size database for a database subsystem. Disk protection is not included in these sizings but can be added to increase the stability of the database.

Increasing the size of your database (DB2 for Linux, UNIX, and Windows only)

DB2 for Linux, UNIX, and Windows Workgroup Edition has a default table size limit of 64 GB with a page size of 4 KB. To increase the capacity of your DB2 for Linux, UNIX, and Windows database, you can create a new table space, IBMDEFAULTGROUP, and choose a larger page size (up to 32 KB). This increases the capacity of the database up to 512 GB per table.

About this task

The following example creates the IBMDEFAULTGROUP table space with a page size of 16 K. This increases the table size capacity to 256 GB.

```
CREATE REGULAR TABLESPACE IBMDEFAULTGROUP IN DATABASE PARTITION GROUP IBMCATGROUP
PAGE SIZE 16384 MANAGED BY DATABASE
USING (FILE 'E:\DB2\NODE0000\SQL00001\IBMDEFAULTGROUP.001' 1500000)
EXTENT SIZE 32
PREFETCH SIZE AUTOMATIC
BUFFER POOL IBM16KBP
OVERHEAD 12.670000
TRANSFER RATE 0.180000
DROPPED TABLE RECOVERY ON
```

If 512 GB of space per table is not enough for your environment, move to DB2 for Linux, UNIX, and Windows Enterprise Edition and using physical or logical partitioning.

The following steps outline the process for using database partitioning with DB2 for Linux, UNIX, and Windows Enterprise Edition:

Procedure

1. Add a new database partition to the DB2 for Linux, UNIX, and Windows instance by running the **db2nrcrt** command.
2. Use the ALTER TABLE statement to add a partitioning key to the tables that you want to partition. For example:

```
ALTER TABLE "CANDLE"."NT_System" ADD PARTITIONING KEY ("Server_Name")  
USING HASHING
```
3. Use the ALTER DATABASE PARTITION GROUP statement to assign the new partition to the database partition group. You can do this either from the command-line or from the DB2 Control Center.
4. Redistribute the data in the database partition group, using the Redistribute Data Wizard in the DB2 Control Center.

What to do next

For additional information about database partitioning, see the following DB2 for Linux, UNIX, and Windows sources:

- *IBM DB2 Universal Database Administration Guide: Planning*
- *IBM DB2 Universal Database Administration Guide: Implementation*
- DB2 for Linux, UNIX, and Windows Information Center at <http://publib.boulder.ibm.com/infocenter/db2help/index.jsp>.

For additional database performance and tuning information, see the *IBM Tivoli Monitoring Administrator's Guide*.

Planning assumptions

The information in the next four chapters is based on the assumptions in this section.

- You have completed the preliminary planning required to determine the size and topology of your environment and your warehousing needs. You may have already installed a monitoring server, a portal server, and some monitoring agents, but you have not yet created the Tivoli Data Warehouse.
- You are not installing IBM Tivoli Monitoring on one computer.
- You will create the Tivoli Data Warehouse database on a different computer from the Tivoli Enterprise Portal Server.
- If installing the warehouse database on Microsoft SQL Server, you will also install the Tivoli Enterprise Portal Server on a Windows-based computer. This restriction applies even if the warehouse database and portal server are installed on separate computers. For example, a portal server on Linux does not support a warehouse database using Microsoft SQL Server.
- If you are upgrading from IBM Tivoli Monitoring V6.1, you have performed the necessary agent warehouse database upgrade steps.

The following sections discuss these assumptions in more detail.

Preliminary planning is complete

The information in the next four chapters assume that your preliminary planning is complete.

For guidance on planning the size of your environment and determining the warehousing capacity you need, see “Sizing your Tivoli Monitoring hardware” on page 34 and “Planning considerations for the Tivoli Data Warehouse” on page 510. The deployment scenarios illustrate your options for where to locate the warehousing components, whether to have multiple Warehouse Proxy Agents reporting to the same hub monitoring server, and whether to deploy single or multiple hub installations.

Environment assumptions

The information in the current chapter and the next four chapters assume an environment with more than one computer.

In environments where monitoring and warehousing needs are minimal, such as a test environment, you can install all IBM Tivoli Monitoring components on one computer: a monitoring server, portal server, data warehouse, Warehouse Proxy Agent, and Summarization and Pruning Agent. If you install all components on a Windows system with either DB2 for Linux, UNIX, and Windows or Microsoft SQL Server, many of the warehouse configuration tasks are automated.

If you want to deploy IBM Tivoli Monitoring on one Windows computer with either of these database platforms, follow the instructions in Chapter 9, “Installing IBM Tivoli Monitoring on one computer,” on page 171.

The data warehouse is remote from the portal server

The procedures described in the next four chapters are based on the assumption that you will install the Tivoli Data Warehouse database remotely from the portal server.

There are two databases involved in IBM Tivoli Monitoring:

- The *Tivoli Enterprise Portal Server database* (or *portal server database*) stores user data and information required for graphical presentation on the user interface. The portal server database is created automatically during configuration of the portal server. It is always located on the same computer as the portal server.
- The *Tivoli Data Warehouse database* (also called the *warehouse database* or *data warehouse*) stores historical data for presentation in historical data views. While it is possible for the warehouse database to be located on the portal server (as in the single computer deployment), it is best to create the database on a remote computer to handle the warehousing needs of most production environments.

No assumption is made about where you will install the Warehouse Proxy Agents and Summarization and Pruning Agent. Either of these agents may be installed on the same computer as the Tivoli Data Warehouse or on a different computer:

- Installing the Warehouse Proxy Agent and Summarization and Pruning Agent on the same computer as the data warehouse minimizes server deployments, simplifies configuration, and eliminates network transmission overhead.
- If you install these components on different computers, ensure that you have a high-speed network connection for best performance.
- In environments with more than one Warehouse Proxy Agent reporting to the same hub monitoring server, install each Warehouse Proxy on a separate computer.

Agent warehouse database upgrade

Some of the monitoring agents have made changes to the warehouse tables that require performing upgrade procedures from IBM Tivoli Monitoring V6.1 before running the V6.3 Warehouse Proxy and Summarization and Pruning Agents.

See “Upgrading the warehouse from IBM Tivoli Monitoring V6.1.x levels” on page 156 for more information.

Firewall considerations for the Tivoli Data Warehouse

Firewalls can be a significant factor in most Tivoli Data Warehouse implementations.

The following minimum requirements must be met:

- The Warehouse Proxy Agent must be able to communicate with:
 - The database
 - The hub monitoring server
 - The agents that are exporting data
- The Summarization and Pruning Agent must be able to communicate with:
 - The database
 - The monitoring server it is configured to communicate with (not necessarily the hub)
 - The Tivoli Enterprise Portal Server

Note: The default Tivoli Enterprise Portal Server interface port of 15001 is also used after the initial connection of the Summarization and Pruning Agent to the portal server over port 1920. Any firewalls between the two need to allow communications on either 15001 or whichever port is defined for any new Tivoli Enterprise Portal Server interface used per the instructions in “Defining a Tivoli Enterprise Portal Server interface on Windows” on page 876.

- The agents exporting data to the Warehouse Proxy Agent must be able to communicate with the Warehouse Proxy Agent.

Appendix C, “Firewalls,” on page 861 contains information on the firewall options available. The IBM redbook *Tivoli Management Services Warehouse and Reporting*, available from <http://www.redbooks.ibm.com>, discusses firewall considerations specific to the Tivoli Data Warehouse in detail.

Compressing historical data transmission

The data from the agent or the monitoring server can be compressed before it is transmitted to the Warehouse Proxy Agent. The data is then decompressed at the Warehouse Proxy Agent before it is inserted in the warehouse database, allowing for efficient usage of network bandwidth between the historical client and the Warehouse Proxy Server.

The following reductions can occur:

- A reduction in timeout errors due to a low bandwidth between clients and Warehouse Proxy Server. Timeout errors on the client side occur when it takes more than 15 minutes for an export of data to complete and be inserted in the warehouse database.
- A reduction in bandwidth usage. Uploading historical data to the warehouse can tie up a significant amount of bandwidth, and also impact other critical applications sharing the network link.
- A reduction in the elapsed time for the upload over slow links.

Use this functionality when a slow network is used or when the network is shared with other critical applications. However, compressing the data does increase CPU usage, both at the agent where it is uploaded and at the Warehouse Proxy Agent where it is decompressed. CPU costs differently on a distributed source than on a z/OS source. In a mixed environment, it is possible to enable or disable the compression option (which is a server option), depending on the source operating system of the client. It is also possible to overwrite a server compression option.

Configuration parameters for the Warehouse Proxy Agent are set in the following files according to operating system:

Windows

`ITM_HOME\TMAITM6\khdenv`

For example: C:\IBM\ITM\TMAITM6\khdenv

Linux

UNIX

ITM_HOME/config/hd.ini

For example: /opt/IBM/ITM/config/hd.ini

The Warehouse Proxy Agent configuration file now contains two new variables that you can modify if necessary:

KHD_SERVER_DIST_COMPRESSION_ENABLE

If this variable is set to Y, then the Warehouse Proxy server allows distributed clients to send compressed data. This variable is set to Y by default.

KHD_SERVER_Z_COMPRESSION_ENABLE

If this variable is set to Y, then the Warehouse Proxy server allows z/OS clients to send compressed data. This variable is set to N by default.

The agent configuration file, if historical data stored at the Tivoli Enterprise Monitoring Agent, or the monitoring server configuration file, and if historical data is stored at the monitoring server, accepts a new warehouse variable:

KHD_CLIENT_COMPRESSION_ENABLE

If set to N, and even if the Warehouse Proxy Agent server has allowed the compression, the historical data is not compressed on this source. This variable does not exist by default.

The new Warehouse Proxy Agent variables are visible in a view in the configuration workspace of the Warehouse Proxy Agent. They also exist on the configuration GUI of the Warehouse Proxy Agent.

You can see if the compression is occurring by using a tracing level to see the data compressed at the client and uncompressed at the server: ERROR (UNIT: khdxmts DETAIL) in the warehouse proxy agent. To check uncompression in the server log file, see if the file contains the following text:

```
khdxmts.cpp,154,"KHD_SendBuff_from_xmit_rep") Uncompressing the data
```

Use ERROR (UNIT: khdxdacl DETAIL) to see the message about compression in the client. To check that the client has overwritten the server variable to not compress the data, see if the file contains the following text:

```
khdxmts.cpp,167,"KHD_SendBuff_from_xmit_rep" Client did not compress the data
```

If compression has failed, you will see the following text in the Warehouse Proxy Server log file, khdxrpcs.cpp:

```
"Client compression failed, data was sent uncompressed"
```

If compression has failed, you will see the following text in the client log file, khdxdacl.cpp:

```
"Warehouse client unable to compress the data, data sent uncompressed"
```

If uncompression has failed, you will see the following text in the Warehouse Proxy Server log file, khdxrpcs.cpp:

```
"Uncompress error"
```

If uncompression has failed, you will see the following text in the client log file, khdxdacl.cpp:

```
"Warehouse client unable to compress the data, data sent uncompressed"
```

Complete the following steps to setup compression:

1. Install or upgrade to the newest release.

2. Configure the Warehouse Proxy Agent, using the GUI. The GUI can pop up during install time, or can be launched by right clicking on the warehouse Proxy entry in the manage Tivoli Monitoring Services panel. The new server compression variables are part of the GUI and can be modified as needed.
3. Compression variables can also be reconfigured using the command-line interface (`itmcmd config hd`).
4. Compression variables can also be reconfigured in the configuration file itself. Configuration parameters for the Warehouse Proxy Agent are set in the following files according to operating system:

Windows

`ITM_HOME\TMAITM6\khdenv`

For example: `C:\IBM\ITM\TMAITM6\khdenv`

Linux

UNIX

`ITM_HOME/config/hd.ini`

For example: `/opt/IBM/ITM/config/hd.ini`

In mixed environments with a previous version of the client and newer version of the server, or new version of the client and a previous version of the server, the compression should not happen.

Next steps

After you have completed your preliminary planning, you are ready to implement your Tivoli Data Warehouse solution.

- Review the *Summary of supported operating systems* section in this chapter to understand your options for operating system platforms, database platforms, and communications between warehousing components, all summarized in a single composite graphic. This section also describes the relationships between warehousing components.
- The next four chapters are organized by database platform. Follow the instructions in the chapter that pertains to the database platform you are using for the Tivoli Data Warehouse database: IBM DB2 for Linux, UNIX, and Windows, IBM DB2 on z/OS, Microsoft SQL Server, or Oracle.

Summary of supported operating systems

This section illustrates the supported operating system platforms for the various warehousing components, the supported database products, and the connections between components.

For more specific information about supported operating systems and database products, including product names and versions, see Chapter 3, “Hardware and software requirements,” on page 75. In this diagram, UNIX refers to any UNIX platform supported by the RDBMS.

Note: The diagram is not meant to suggest that each warehousing component must be installed on a separate computer. Multiple components can exist on the same computer provided there is operating system support for all installed components. Connections between components must be configured whether or not they exist on the same computer. (Although the diagram makes no assumptions about where components are installed, the procedures described in the next four chapters assume a Tivoli Data Warehouse that is remote from the portal server.)

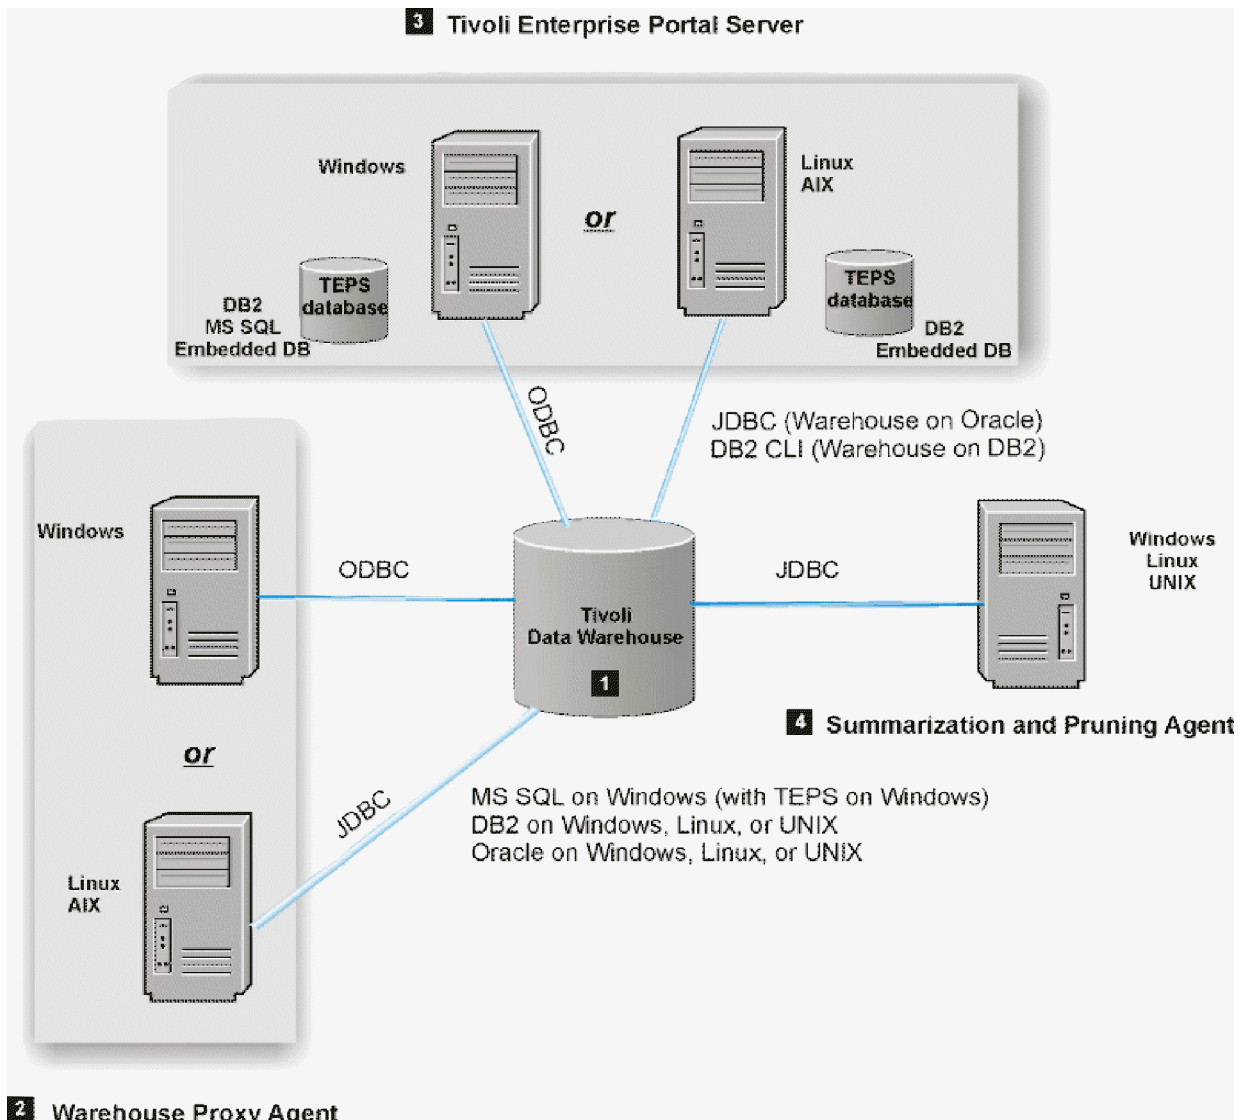


Figure 96. Summary of support for the Tivoli Data Warehouse

In the following discussion, numbered product components correspond to the numbers on the diagram.

1 Tivoli Data Warehouse database

Data collected by monitoring agents is stored at intervals in short term history files. The data in the short term history files is referred to as *historical data*. The short term history files are located either at the agents or at the monitoring server (hub or remote) to which the agents report. (An administrator can determine where the data is stored. The monitoring agents and monitoring server are not shown in the diagram.)

The historical data is sent from its temporary storage location (at the monitoring agents or monitoring server) to the Warehouse Proxy Agent at a preset interval (either every hour or every 24 hours). The Warehouse Proxy Agent inserts the data it receives to the Tivoli Data Warehouse.

Data in the short term history files that is older than 24 hours is pruned when the monitoring agent or monitoring server receives an acknowledgment that the data has been successfully inserted into the Tivoli Data Warehouse. (The pruning of binary files is not the pruning performed by the Summarization and

Pruning Agent. The Summarization and Pruning Agent prunes data in the Tivoli Data Warehouse.) The result of these operations is that, at any given time, the short-term historical files contain data that is less than 24 hours old, and the Tivoli Data Warehouse contains long-term historical data that is older than 24 hours.

The Tivoli Data Warehouse database can be created using Microsoft SQL Server, IBM DB2 Database for Linux, UNIX, and Windows, Oracle, or IBM DB2 on Z/OS on the indicated operating system platforms. Note that the warehouse database is supported on Microsoft SQL Server only if the Tivoli Enterprise Portal Server is installed on Windows. This condition applies even if the warehouse database and portal server are installed on separate computers. For example, a portal server on Linux does not support a warehouse database on Microsoft SQL Server.

2 Warehouse Proxy Agent

A Warehouse Proxy Agent running on Windows uses an **ODBC** connection to send the collected data to the warehouse database. On a 32-bit Windows system, this requires a 32-bit ODBC driver, but on 64-bit Windows, it requires a 64-bit ODBC driver by default (although sites running a 32-bit Warehouse Proxy Agent under 64-bit Windows can configure a 32-bit ODBC driver for the agent to use; the necessary instructions are provided in this manual). The type of ODBC driver your site uses is determined not by the operating system but rather by the type of Warehouse Proxy Agent your site runs: the 32-bit agent requires a 32-bit ODBC driver, whereas the native 64-bit agent requires a 64-bit ODBC driver.

A Warehouse Proxy Agent running on Linux or AIX uses a **JDBC** connection.

The Warehouse Proxy Agent will do the following as needed:

- Create new tables and indexes.
- Alter existing tables.

For example, add a new column to a table. This can occur if the table was created for an older version of an agent but data is being exported for the first time from a newer version of the agent that has more attributes.

Firewall considerations:

- The Warehouse Proxy Agent must be able to communicate with the hub monitoring server to register its network address. This address will then be replicated to all remote monitoring servers. If communication is impossible for firewall reasons, a variable indicating the Warehouse Proxy Agent network addresses can be set on the hub monitoring server and remote monitoring servers: `KPX_WAREHOUSE_LOCATION`.
- The Warehouse Proxy Agent must be able to communicate with the RDBMS.
- The agents exporting data must be able to communicate with the Warehouse Proxy, *and* the Warehouse Proxy Agent must be able to communicate with the agents (in other words, there must be two-way communication between agents and the Warehouse Proxy Agent).

Appropriate ports must be open to support these communication channels (see Appendix C, “Firewalls,” on page 861).

Depending on the database system you're using with your Tivoli Data Warehouse, make sure you complete one of these steps when installing the Warehouse Proxy Agent:

- If your site hosts its Tivoli Data Warehouse on DB2 Database for Linux, UNIX, and Windows, complete this step: “Configuring an ODBC data source for a DB2 data warehouse” on page 551.
- If your site hosts its Tivoli Data Warehouse on Microsoft SQL Server, complete this step: “Configuring an ODBC data source for a Microsoft SQL data warehouse” on page 602.
- If your site hosts its Tivoli Data Warehouse on Oracle, complete this step: “Configuring an ODBC data source for an Oracle data warehouse” on page 627.

Note: If your site has previously installed the 32-bit agent and you wish to upgrade to the 64-bit agent, you must first uninstall the 32-bit agent; see “Uninstalling the Warehouse Proxy” on page 937. This might also require removing the 32-bit ODBC driver you configured for the Warehouse Proxy Agent; see also “Removing the ODBC data source connection” on page 938.

When uninstalling a 32-bit agent so your site can upgrade to a 64-bit agent, remember that any customizations your site has made to the KHDENV configuration file must be redone. Also, a 64-bit ODBC driver must be installed for the 64-bit agent to use.

Usual process for activating the 64-bit Warehouse Proxy Agent:

1. Install the ODBC 64-bit data source, if not installed already.
2. Install the 64-bit agent.
3. Configure the agent to use a 64-bit data source.
4. Start the agent.

To determine which Warehouse Proxy Agent is running (32-bit or 64-bit), open the Windows Task Manager. The 32-bit agent has “*32” next to the process name.

3 Tivoli Enterprise Portal Server

The Tivoli Enterprise Portal Server retrieves historical data for display in historical data views in the Tivoli Enterprise Portal and in monitoring dashboard applications for the Dashboard Application Services Hub. It retrieves short-term historical data from the binary files on the monitoring agents or monitoring server. It retrieves long-term historical data from the Tivoli Data Warehouse.

In Figure 96 on page 520, the Tivoli Enterprise Portal Server is shown with the *portal server database* (designated as *TEPS database* in the diagram). The portal server database stores user data and information required for graphical presentation on the user interface. Before you install and configure the portal server, you must install the database platform (RDBMS) to be used for the portal server database (DB2 for Linux, UNIX, and Windows or Microsoft SQL Server) on the same computer. The portal server database is created automatically during configuration of the portal server.

Although the portal server database is not considered a warehousing component, it is included in the diagrams in this and the following chapters because it can affect the installation and configuration tasks required for the Tivoli Data Warehouse database. For example, the database client already installed for the portal server database can connect to a remote warehouse database, provided both databases use the same database platform. There is no need to manually install another client.

If the portal server is installed on Windows, it uses an ODBC connection to request and retrieve historical data from the warehouse database. If the portal server is installed on Linux or AIX, it communicates with the warehouse database through a JDBC connection, if the warehouse is installed on Oracle, or through a proprietary DB2 CLI connection if the warehouse is installed on DB2 for Linux, UNIX, and Windows.

4 Summarization and Pruning Agent

The Summarization and Pruning Agent retrieves detailed data from the warehouse database, aggregates or prunes the data, and returns the processed data to the warehouse. Communication takes place through a JDBC connection, regardless of the operating system on which the Summarization and Pruning Agent is installed.

The Summarization and Pruning Agent will create tables, indexes, and views as needed. This could happen in either of the following situations:

- Summarization is enabled for the first time for an attribute group.
- Additional summarizations are enabled for an attribute group.

In addition, the Summarization and Pruning Agent, like the Warehouse Proxy Agent, may alter existing tables by adding new columns.

Firewall considerations:

- The Summarization and Pruning Agent must be able to communicate with the RDBMS.
- The Summarization and Pruning Agent must be able to communicate with the Tivoli Enterprise Portal Server.
- The Summarization and Pruning Agent must be able to communicate with the Tivoli Enterprise Monitoring Server it is configured to use.

Chapter 20. Schema Publication Tool

The schema publication tool is an SQL editor that allows you to generate SQL statements and create your database objects.

With the schema publication tool you can perform the following tasks:

- Generate the SQL statements needed to create the database objects required for initial setup of the Tivoli Data Warehouse.
- Create the necessary database objects whenever either historical collection is enabled for additional attribute groups or additional summarizations are enabled. This is referred to as running the schema publication tool in *updated* mode. Updated mode is also used if agents are upgraded to a new release. An upgrade might add new columns to existing tables and the schema publication tool will generate the appropriate alter table statements.

Generating SQL for data warehouse tables

You can use the schema publication tool to generate the SQL statements needed to create the database objects required for initial setup of the Tivoli Data Warehouse.

Before you begin

The schema publication tool is installed with the Summarization and Pruning agent. You should perform this task after product installation and after configuring the Summarization and Pruning agent, but before starting the Warehouse Proxy agent and the Summarization and Pruning agent for the first time.




About this task

By default, the database objects required for the data warehouse are created automatically by the installer, the Warehouse Proxy agent and the Summarization and Pruning agent. The schema publication tool enables you to create the data warehouse database objects manually rather than allowing them to be created automatically. There are several situations in which you might want to do this:

- You might want to use a separate database administration user ID for creating the tables, rather than granting permission to the Tivoli Data Warehouse user ID specified during installation.
- You might want to customize the SQL before creating the tables in order to accommodate performance considerations, security policies, or other issues unique to your environment.

The schema publication tool is a script that generates the SQL required to create the data warehouse tables, indexes, functions, views, and ID table inserts required for the selected products. You can then modify the generated SQL files before using them to create the tables and indexes.

Procedure

1. Create a new response file by making a copy of the sample response file:
 -  `itm_install_dir\TMAITM6\tdwschema.rsp`
 -   `itm_install_dir/arch/bin/tdwschema.rsp`
2. Using an ASCII text editor, edit the response file to indicate the options you want to use. The keywords in the response file affect which SQL statements are generated, as well as other options:

KSY_PRODUCT_SELECT=category

The category of products for which you want to generate SQL files:

Value	Description
installed	All installed Tivoli Enterprise Portal Server products. Note: This can produce DDL for a large number of tables.
configured	All configured portal server products. Products that have historical collections defined. DDL is only generated for the attribute groups that have historical collections defined. DDL for summary tables is generated based on the summarization and pruning settings. For example, if an attribute group has hourly and monthly summarization enabled, DDL is generated for only the hourly and monthly summary tables.
updated	Configured portal server products with configuration changes that have not yet been deployed to the database. The updated value captures the changes that are necessary to bring the Tivoli Data Warehouse up to date to the current configuration. If you require the entire schema that is configured, you must use the configured value. Examples of updates: <ul style="list-style-type: none"> • Historical collections created for an attribute group that previously had no collections. • Enabling additional summarizations for one or more attribute groups. • Software update that changes existing attribute groups. Existing warehouse tables will need to have columns added. • Database compression is enabled at the database and the Summarization and Pruning agent is configured to use database compression. DDL is generated to compress existing tables and indexes.

This keyword is required.

KSY_PRODUCT_FILTER=*product_filter*

An optional filter to indicate that only certain specific products are included. (If you do not specify a filter, all products in the specified category are included by default.) Specify the three-letter product codes of the products you want to include, separated by commas. You can find these codes by using the **tacmd histListProduct** command (for more information, refer to the *IBM Tivoli Monitoring Command Reference*).

KSY_TABLE_FILTER=*list_of_tables*

You can use this filter in addition to KSY_PRODUCT_FILTER. Use the following command to get the list of tables that are available for a given product:

```
tacmd histListAttributeGroups -t <productcode>
```

Replace each space in the attribute group name with an underscore character.

KSY_SUMMARIZATION_SELECTION=*summarization_filter*

An optional filter to indicate that only certain summarization options are to be included in the generated tables:

Value	Description
H	Hourly
D	Daily

Value	Description
W	Weekly
M	Monthly
Q	Quarterly
Y	Yearly

This keyword is valid only with KSY_PRODUCT_SELECT=installed.

KSY_SQL_OUTPUT_FILE_PATH=*path*

An optional path to the directory where the generated SQL files are to be written. If you do not include this keyword, the current working directory is used.

For more details and the complete syntax for each keyword, refer to the comments in the tdwschema.rsp sample response file.

3. Make sure the Tivoli Enterprise Portal Server is started.
4. Run the schema publication tool script using the appropriate syntax for your operating system.

- **Windows**

```
tdwschema -rspfile response_file
```

- **Linux**

- **UNIX**

```
tdwschema.sh -rspfile response_file
```

The SQL files for the products specified in the response file are generated and written to the directory indicated by the KSY_SQL_OUTPUT_FILE_PATH keyword (or to the current working directory, if no output directory is specified).

5. Make any necessary changes to the generated SQL files. For example, you might want to partition tables or assign tables to table spaces.

Note: Do not change the names of any tables specified in the generated SQL files.

6. Use the appropriate tools to run the SQL queries to create the warehouse tables, indexes, views, inserts, and functions for your relational database. Execute the scripts in this order:
 - a. tdw_schema_table.sql
 - b. tdw_schema_index.sql
 - c. tdw_schema_view.sql
 - d. tdw_schema_insert.sql
 - e. tdw_schema_function.sql

The following examples are for the DB2 commands:

```
db2 -tvf tdw_schema_table.sql
db2 -tvf tdw_schema_index.sql
db2 -tvf tdw_schema_view.sql
db2 -tvf tdw_schema_insert.sql
db2 -td# -f tdw_schema_function.sql
```

Note: The different invocation for the tdw_schema_function.sql.

Using the schema publication tool in updated mode

Use the schema publication tool's updated mode to create the necessary database objects whenever either historical collection is enabled for additional attribute groups or additional summarizations are enabled.

About this task

Note: If you run the schema publication tool in updated mode on an existing IBM Tivoli Monitoring version 6.1 Tivoli Data Warehouse database, the following will not be done:

- In an IBM Tivoli Monitoring V6.1 Tivoli Data Warehouse database, indexes on aggregate tables for a given attribute group do not include all of the key columns for the attribute group. This causes

performance problems with the Summarization and Pruning agent. The updated mode does not generate SQL to recreate the indexes to optimize performance.

- In IBM Tivoli Monitoring V6.1, all character data was stored in fixed-length CHAR columns. In IBM Tivoli Monitoring V6.2 and 6.3, VARCHAR is used which greatly reduces disk requirements and improves performance. The updated mode does not generate SQL to convert CHAR columns to VARCHAR.
- In IBM Tivoli Monitoring V6.1, all columns allowed NULL values. In IBM Tivoli Monitoring V6.2 and V6.3 some columns that can never be NULL were changed to include a NOT NULL constraint. In large tables this can save significant disk space. For DB2, this greatly reduced the disk requirements for indexes. The updated mode does not generate SQL to set columns to NOT NULL.
- If using DB2 for Linux, UNIX, and Windows and if the IBM Tivoli Monitoring V6.1 database is not migrated properly, the schema tool may produce SQL that fails. The tool may generate ALTER TABLE statements that cause a table not to fit into the table's page size.

Procedure

1. Create a new response file by copying the sample response file:
 - **Windows** Copy `itm_install_dir\TMAITM6\tdwschema.rsp`.
 - **Linux** **UNIX** Copy `itm_install_dir/arch/bin/tdwschema.rsp`.
2. Using an ASCII text editor, edit the response file as follows: **KSY_PRODUCT_SELECT=updated**. (See the description above.)

You may also specify an output path via the **KSY_SQL_OUTPUT_FILE_PATH=***path* parameter, as explained above.
3. Using either the historical configuration interface within the Tivoli Enterprise Portal or the historical configuration CLI, make the desired changes to your site's historical configuration. If enabling historical collection for new attribute groups, configure but do not start historical collection. If collection is started, the warehouse proxy agent may attempt to create the database objects before you have a chance to generate, edit, and execute the SQL.
4. Ensure the Tivoli Enterprise Portal Server is started.
5. Run the schema publication tool script:

- **Windows**
`tdwschema -rspfile response_file`
- **Linux** **UNIX**
`tdwschema.sh -rspfile response_file`

where *response_file* is the name of the response file you edited in step 1. The SQL files for the products specified in the response file are generated and written to the directory indicated by the **KSY_SQL_OUTPUT_FILE_PATH** keyword (or to the current working directory, if no output directory is specified).

6. Make any necessary changes to the generated SQL files. For example, you might want to partition tables or assign tables to table spaces.

Note: Do not change any of the following in the generated SQL files:

- Table or view names
 - Table column names
 - Table column sizes or data types
7. Use the appropriate tools to run the SQL queries to create the warehouse tables, indexes, views, inserts, and functions for your relational database. Run the scripts as the Tivoli Data Warehouse user in this order:
 - a. `tdw_schema_table.sql`
 - b. `tdw_schema_index.sql`
 - c. `tdw_schema_view.sql`
 - d. `tdw_schema_insert.sql`

- e. `tdw_schema_function.sql`
- 8. Using either the historical configuration interface within the portal or the historical configuration CLI, start historical collection for the newly configured attribute groups.

Using the schema publication tool in migrate mode

You can use the schema publication tool to migrate existing non-partitioned tables to partitioned tables.

For detailed information on migrating non-partitioned tables, see “Tivoli Data Warehouse range partition migrations” in the *IBM Tivoli Monitoring Administrator's Guide* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/adminuse/itm_admin.htm).

Using the Schema Publication Tool for database compression

The Schema Publication Tool allows you to generate DDL scripts that can be used against a database capable of compression.

The DDL scripts alter the tables and indexes, including the log control tables. Executing those scripts will help you to reduce the amount of disk space used and also improve database performance. Reducing storage subsystem costs can result in substantial cost savings. The Schema Publication Tool generated scripts can be modified to accommodate local security concerns, point to a particular tablespace, or otherwise customize the generated DDL.

Note:

1. The Summarization and Pruning Agent must be configured and have an available database connection.
2. At this time, compression on DB2 z/OS is not supported.
3. Compression increases CPU use by the database application.
4. When using DB V9.1 a reorganization is required for each table and index created with compression enabled, whether it is existing or new data. The reorganization creates the required compression dictionary.
5. DB2 V9.5 and greater automatically creates the required compression dictionary. A reorganization is required to compress existing data, otherwise only new data is compressed.
6. The Schema Publication tool uses only the compression setting stored in the Summarization and Pruning configuration file for all tables and indexes. This setting includes tables normally generated by the Warehouse Proxy Agent such as raw tables and associated control tables.
7. If the compression setting is turned off after table and indexes are compressed, the Schema Publication Tool will not generate the DDL to change them back to uncompressed mode.

Ensure a compression enabled database application server is installed and properly licensed. Summarization and Pruning Agent must be configured and an active database connection available. When using Oracle, compression is on a block basis (while on DB2 it is on the entire table). A larger block size could result in significantly larger compression. Block size is defined at database creation and cannot be changed, and block size needs to be a multiple of the underlying operating system block size.

The amount of storage reduction depends on a number of factors including the following consideration: the structure of the data makes an enormous difference. More compression is generated if there are repeated values in the data since all of the database applications use dictionary compression algorithms. The total amount of storage saved is proportional to the size of the table. However, the compression ratio is based on the data itself.

The compression settings and any warnings can be found in the Summarization and Pruning Agent log files and in the Warehouse Proxy Agent and Summarization and Pruning Agent's Tivoli Enterprise Portal configuration workspaces. If the Schema Publication Tool finds the DBMS does not support compression

and DB compression is enabled, then the DB compression will be set to NO, a log entry made, and also the script will issue a message to that effect. All DDL will then be created without compression.

Complete the following steps to enable compression:

1. In the Summarization and Pruning Agent environment file, set the KSY_DB_COMPRESSION environment variable to Y. The default is N.
For Windows systems, the environment file name and location is *install dir\TMAINT6\KSYENV*. For UNIX systems, the environment file name and location is *install dir/config/sys.ini*.
2. Update the required settings in the Schema Publication Response file and run the *tdwschema.bat* (Windows) or *tdwschema.sh* (UNIX) command. The response file property for this setting is: KSY_PRODUCT_SELECT = "updated".

Note: In the following situations there will be separate values and potentially conflicting options for the compression option:

- The Warehouse Proxy Agent and the Summarization and Pruning Agent have separate compression options and can be out of sync.
- If there are multiple Warehouse Proxy Agents, the values could be out of sync.

In all cases, the Schema Publication Tool uses only the Summarization and Pruning Agent setting. No mechanism is provided to coordinate the settings.

Chapter 21. Tivoli Data Warehouse range partitioning

Range partitioning is a database data organization feature that allows for quicker pruning of data from the database in a single operation. Range partitioning permits the fast roll out of data without having to perform a resource intensive DELETE operation on blocks of rows. Range partitioning also improves query performance when the partitioning key is part of the query clause.

Note: Supported databases for Tivoli Data Warehouse database range partitioning:

- IBM DB2 Enterprise Edition 9.7 or higher.
- DB2 for z/OS 9.1 or higher.
- Oracle 11g Release 1 or higher, with the partitioning option.

Performance is improved because the database can discard unnecessary partitions from the list of blocks that must be fetched from disk. With table partitioning you can significantly improve pruning and query performance in large Tivoli Data Warehouse databases. As of Tivoli Monitoring version 6.3, Tivoli Data Warehouse range partitioning provides the following functionality:

- You can configure the Warehouse Proxy Agent and the Summarization and Pruning Agent to create partitioned tables.
- You can use the schema publication tool to generate the DDL for partitioned tables, for a database administrator to run later.
- The Summarization and Pruning Agent is responsible for managing partitions based on your configuration, that is adding partitions or pruning partitions.
- Depending on your Tivoli Data Warehouse database, the Summarization and Pruning Agent prunes data from partitioned tables as follows:
 - On DB2 for Linux, UNIX, and Windows or Oracle: The Summarization and Pruning Agent prunes data from partitioned tables by dropping the appropriate partitions.
 - On DB2 on z/OS: The Summarization and Pruning Agent prunes data by rotating partitions.

If the table being pruned is not partitioned or is partitioned using a different scheme than the one used by Tivoli Data Warehouse, rows are pruned using SQL DELETE (the existing pruning logic). It is recommended best practice to use the Warehouse Proxy Agent and the Summarization and Pruning Agent to enable partitioning.

Range partitioning using DB2 for Linux, UNIX, and Windows or Oracle: The partition range used by Tivoli Data Warehouse is one day and the partition is named `PYYYYMMDD`. A catch all partition with an additional suffix of `_MV` is also created and will contain any data older than the day that the table was created by either the Warehouse Proxy Agent (detail tables) or the Summarization and Pruning Agent (summary tables). The data in the catch all partition is not pruned until the entire partition is outside of the retention period of the table. Therefore, when data is deleted from a partitioned table by removing or detaching partitions, data older than the configured retention period can still exist in the table after pruning is performed. The partitioning key is one column: the aggregation timestamp, which is usually called `WRITETIME`.

Range partitioning using DB2 on z/OS: The partition range used by Tivoli Data Warehouse is one day and the partition is named using an incremental number beginning with 1. Partitions are rotated, they cannot be detached from a table. Rotating partitions allow old data to roll off, while reusing the partition for new data. In order to control the number of partitions added, the Summarization and Pruning Agent does not allow the addition of partitions when the maximum number is reached. The maximum partition number is equal to the sum of 1 (the partition holding the current day's data) plus the number of retention days, weeks, months, quarters, or years. This retention period depends on the pruning settings for the aggregated table and the number of upward partitions set in the `KSY_PARTITIONS_UPWARD`

variable. It is recommended to set some pruning parameters for a table to avoid the unlimited growth of partitions against a DB2 on z/OS Tivoli Data Warehouse database.

The environment variables listed in this section can be configured from the agent configuration GUI, the Tivoli Enterprise Portal, and the Command-Line Interface. Errors are reported via the appropriate Tivoli Data Warehouse agent and available for situation processing and workspaces.

When tables are pruned by dropping partitions, the reported statistics for the number of rows deleted is an estimate only. The estimate is based on the number of rows that were in the partition the last time the database statistics were updated for the table. The database statistics must be updated on a regular basis so that this number is as accurate as possible and to help improve the performance of SQL against the table.

Note: Data partitioning might require additional licenses for your database management system. Consult with your database administrator for any rules or restrictions imposed by your database implementation before using range partitioning.

Specifying range partitioned tables for the Warehouse Proxy Agent

Table partitioning only occurs when new tables are created in the Warehouse Proxy Agent. There is no change to existing tables that were created before Tivoli Monitoring V6.3.

You can configure database table partitioning from the agent configuration GUI, by editing the environment file, or through the Command-Line Interface (for more information, see the *Warehouse Proxy Agent User's Guide*.)

Any export of data for a new attribute group will generate a new partitioned table if partitioning is enabled and if the Tivoli Data Warehouse database supports partitioning.

The screenshot shows the 'Configure Warehouse Proxy' window with the 'Agent Parameters Details' tab selected. On the left, 'Database Type' is checked and 'Agent Parameters' is unchecked. The main area contains the following fields and options:

- *ODBC DSN: ITM Warehouse
- *Username: itmuser
- *Password: [masked]
- *Confirm Password: [masked]
- Test connection button
- Warehouse TEMS List: [empty]
- ☒ Use Batch
- ☐ Database Compression
- ☐ Warehouse Compression for Z/OS Sources
- ☒ Warehouse Compression for Distributed Sources
- ☒ Database Table Partitioning
- *Number of future partitions to maintain: 10
- Default table container: [empty]
- Default index container: [empty]

At the bottom are buttons for Back, Next, Home, OK, and Cancel.

Figure 97. Configure Warehouse Proxy - Database Table Partitioning

Note: Number of future partitions to maintain must contain a value, regardless of whether or not Database Table Partitioning is selected.

Configuration parameters

There are two parameters that must be set in the Warehouse Proxy Agent environment file.

KHD_PARTITION

Controls whether partitioning is enabled (Y) or disabled (N). When disabled, non-partitioned tables and non-partitioned indices are created and the old pruning behavior is in effect. When enabled, partitioned tables and partitioned indices are created (if supported by the database) and partitions are maintained. The default setting is N.

KHD_PARTITIONS_UPWARD

Number of partitions to maintain after the current date. This allows for an outage to occur and not cause disruption on the monitoring system, permitting agents to upload their collected historical data. Minimum value of 1, maximum value of 10. The default setting is 10.

Note: If there are multiple Warehouse Proxy Agents in your Tivoli Monitoring environment you must configure each of them. Use the same partitioning values for all of the Warehouse Proxy Agents. For the KHD_DEFAULT_TABLE_CONTAINER and KHD_DEFAULT_INDEX_CONTAINER, see “Environment variables” on page 538 in Chapter 22, “Specifying where tables and indices are created,” on page 537.

After the variables in the Warehouse Proxy Agent are set, run the Warehouse Proxy Agent. Schedule the Summarization and Pruning Agent to run at least every day, since only the Summarization and Pruning Agent manages the addition to, and deletion of (or rotation of), partitions from a table. See “Specifying range partitioned tables for the Summarization and Pruning Agent” to understand how the Summarization and Pruning Agent manages the partitions created by the Warehouse Proxy Agent and the partitions created by the Summarization and Pruning Agent. New detailed partitioned tables are created when the Warehouse Proxy Agent processes exports for an attribute group for the first time.

Errors are reported via the appropriate Tivoli Data Warehouse agent and available for situation processing and workspaces.

Specifying range partitioned tables for the Summarization and Pruning Agent

Table partitioning only occurs when new tables are created in the Summarization and Pruning Agent. There is no change to existing tables that were created before Tivoli Monitoring V6.3.

You can configure database table partitioning from the agent configuration GUI, by editing the environment file, or through the Command-Line Interface (for more information, see the *Warehouse Summarization and Pruning Agent User's Guide*.)

When partitioning is enabled, the Summarization and Pruning Agent performs the following functions:

- Creates new summary tables as partitioned tables.
- Adds partitions to existing partitioned tables (including the detail tables created by the warehouse proxy) based on your configuration.
- Prunes data from partitioned tables by dropping the appropriate partitions.

If you partition a table without following the Tivoli Data Warehouse partition naming scheme, the Summarization and Pruning Agent will not take advantage of the partitions and data pruning will use the old method of deleting rows in the table.

Each time the Summarization and Pruning Agent is executed, it looks to see if it needs to create a new summary table. The table is only created if it doesn't already exist. Each partition other than the catch all, only contains one day of data. However, separate partitions are created as follows:

- A partition is created for each day for the detail and hourly summary table.
- For all other summary tables, only one partition is created for the first day of the summary interval:
 - For weekly, one partition is created for the start of the week (weeks can be configured to start on Sunday or Monday).
 - For monthly, one partition is created for the first day of the month.
 - For quarterly, one partition is created for the first day of the quarter.
 - For yearly, one partition is created for the first day of the year.

The catch all partition might contain data older than the pruning window that is configured. The data will not be pruned until the catch all partition is dropped and recreated.

Figure 98. Configure Summarization and Pruning Agent - Database Table Partitioning

Note:

1. **Number of future partitions to maintain** must contain a value, regardless of whether or not **Database Table Partitioning** is selected.
2. If you are using DB2 for Linux, UNIX, and Windows, the DB2 LOCKTIMEOUT must not be set to the default value of -1, which means infinite wait. Adding and detaching partitions require a table lock, which might not be possible to obtain if there is a lot of activity on the table causing locks to be held. Not setting a table lock timeout can cause the Summarization and Pruning Agent to stop processing tables while it is waiting for the table lock.

Configuration parameters

There are three parameters that must be set in the Summarization and Pruning Agent environment file.

KSY_PARTITION

Controls whether partitioning is enabled (Y) or disabled (N). When disabled, non-partitioned tables and non-partitioned indices are created and the old pruning behavior is in effect. When enabled, partitioned tables and partitioned indices are created (if supported by the database) and partitions are maintained. The default setting is N.

KSY_PARTITIONS_UPWARD

Number of partitions to maintain after the current date. This allows for an outage to occur and not cause disruption on the monitoring system, permitting agents to upload their collected historical data. Minimum value of 1, maximum value of 10. The default setting is 10.

KSY_GROUP_BY_THRESHOLD

Number of past days used in the Summarization and Pruning agent discovery query. This is used to improve the Summarization and Pruning agent runtime and to limit the number of partition locks placed on the table. The default setting is 0.

For the KSY_DEFAULT_TABLE_CONTAINER and KSY_DEFAULT_INDEX_CONTAINER, see “Environment variables” on page 538 in Chapter 22, “Specifying where tables and indices are created,” on page 537.

After the variables in the Summarization and Pruning Agent are set, schedule the Summarization and Pruning Agent to run at least every day. Only the Summarization and Pruning Agent is responsible for managing the partitioned tables. Otherwise insertions into tables will fail if there is not a partition defined for the data being inserted.

Partitions are added to the existing partitioned detailed tables, and the existing summarized tables only if the Summarization and Pruning Agent is running and when the upward partition(s) do not exist yet. If the Summarization and Pruning Agent runs every day, a new partition should be added every day at the detail, hourly, and daily level. Weekly aggregated tables will get a new partition once a week, monthly aggregated tables once a month, quarterly aggregated tables once a quarter, and yearly aggregated tables once a year.

Partitions are dropped or detached by the Summarization and Pruning Agent when pruning settings in the Tivoli Enterprise Portal exist against the corresponding attribute group.

Tip: SQL Queries that use the partitioned tables will not require any modification, although queries should always use the partitioning key, usually the WRITETIME table column, to avoid scanning all partitions. Failure to do so can result in locks being acquired which can prevent partitions from being dropped when the Summarization and Pruning Agent is running. Only the following values can be used with the partition key columns:

- Equals
- Not equals
- Less than
- Greater than
- Less than or equals
- Greater than or equals

If a function is used, the database query optimizer will not be able to eliminate partitions.

SQL queries against a DB2 z/OS Tivoli Data Warehouse database should use uncommitted reads via WITH UR to prevent unnecessary locks and failures during the Summarization and Pruning Agent runs.

Errors are treated as any regular error, meaning that the operation being attempted will fail and result in that upload (if Warehouse Proxy Agent) or aggregation (Summarization and Pruning Agent) to fail. Errors are reported via the appropriate Tivoli Data Warehouse agent and available for situation processing and workspaces.

Specifying range partitioned tables for the schema publication tool

As of Tivoli Monitoring V6.3, the schema publication tool allows for the creation of tables with partitions using SQL scripts. This can be used when the user ID for the Warehouse Proxy Agent or Summarization and Pruning Agent does not have the required access permission to create tables.

The schema publication tool is a part of the Summarization and Pruning Agent. Therefore the partitioning settings of the Summarization and Pruning Agent are used when the schema publication tool is executed.

Chapter 22. Specifying where tables and indices are created

For performance purposes a database administrator might impose rules preventing the creation of tables in the default table space. You can specify the tablespaces that are to be used as containers for the Tivoli Data Warehouse tables and indices. The containers are only used when creating new tables. Existing tables are not moved to the specified containers. Tivoli Data Warehouse tables can be placed on separate drives from other tables in the database, indexes can be placed on separate drives from the table data, and so on.

On DB2 for Linux, UNIX, and Windows there are two types of tablespaces: system managed and database managed. For system managed, one or more directories is specified. Data can be added to tables until there is no more room on the file system. For database managed tablespaces, one or more files are created. The space is pre-allocated and data can be added until the files are filled. Database managed tablespaces can perform better than system managed, but they require more maintenance. The database administrator can add more containers as needed, but the table space usage must be monitored.

If only the table tablespace is specified for the Warehouse Proxy Agent or the Summarization and Pruning Agent, the indexes are automatically created in the same table space as the table. The table space can be either a system managed or database managed table space. If a table tablespace and an index table space are specified, both must be database managed tablespaces.

For DB2, tablespaces have a page size of 4K, 8K, 16K, or 32K. Tables must have rows that are less than the page size. However, there is overhead for each row, so the maximum record size is not the same as the page size. On DB2 for Linux, UNIX, and Windows, the row length of a table in the database manager cannot exceed the following sizes:

- 4005 bytes in a table space with a 4K page size
- 8101 bytes in a table space with an 8K page size
- 16293 bytes in a table space with an 16K page size
- 32677 bytes in a table space with an 32K page size

To avoid errors when creating tables, the Warehouse Proxy Agent or the Summarization and Pruning Agent will determine whether or not a table will fit in the specified table space and take one of the following actions:

- If the table does fit in the configured table space, create the table in the table space.
- If the table does not fit in the configured table space, remove both the table and index table space (if an index table space is configured) clauses from the CREATE statement. This causes DB2 to determine the table space that should be used based on the record size of the table (this is the existing Warehouse Proxy Agent and the Summarization and Pruning Agent behavior). The INDEX table space must also be dropped because DB2 may select a system managed table space which would cause the create table to fail.

The schema publication tool works in a similar way. If a table or index table space are specified, the schema tool will (if the tablespaces exist) determine whether or not each table will fit in the table space and take one of the following actions:

- If the table fits, generate CREATE statements that create the table and indices in the specified tablespaces.
- If the table does not fit, drop the table and index table space clauses from the CREATE statement. Warning comments in the generated DDL file, before the CREATE statement, indicate that the clauses were dropped and explain why they were dropped.

If the tablespaces do not exist when the schema publication tool is executed, the table sizes cannot be checked. The schema tool will still include the table and index table space clauses in the generated DDL, but will output a comment at the beginning indicating that the configured tablespaces do not exist. This scenario could happen in a couple of different ways:

- The database administrator creates the tablespaces later just after running the generated DDL.
- The generated DDL is executed on a different DB2 database that does have the specified tablespaces.

Note: You cannot specify where tables and indices are created if you are using DB2 on z/OS or Microsoft SQL. If any of the variables in this section are set using DB2 on z/OS or Microsoft SQL, a warning message in the trace log file states that the default index container or table container is not supported for these database types. The Summarization and Pruning Agent and Warehouse Proxy Agent should execute successfully without taking into account the content of these variables.

Environment variables

There are two variables that must be set in the Warehouse Proxy Agent environment file to specify the table space for the table and the indices.

KHD_DEFAULT_TABLE_CONTAINER

Name of the database container where tables should be placed.

KHD_DEFAULT_INDEX_CONTAINER

Name of the database container where indices should be placed. For DB2, when an index container is specified both the index and table containers must be database managed table spaces.

There are two variables that must be set in the Summarization and Pruning Agent environment file to specify the table space for the table and the indices:

KSY_DEFAULT_TABLE_CONTAINER

Name of the database container where tables should be placed.

KSY_DEFAULT_INDEX_CONTAINER

Name of the database container where indices should be placed. For DB2, when an index container is specified both the index and table containers must be database managed table spaces.

After the variables in the Warehouse Proxy Agent are set, run the Warehouse Proxy Agent. After the variables in the Summarization and Pruning Agent are set, schedule the Summarization and Pruning Agent to run at least every day. You will see the tables and indices created in the table spaces specified in the variables.

Chapter 23. Tivoli Data Warehouse solution using DB2 for Linux, UNIX, and Windows

Use the information and instructions in this chapter to implement a Tivoli Data Warehouse solution using DB2 for Linux, UNIX and Windows for the warehouse database.

The following table lists the goals for creating a DB2 for Linux, UNIX and Windows solution.

Table 81. Goals for creating a Tivoli Data Warehouse solution using DB2 for Linux, UNIX, and Windows

Goal	Where to find information
Review your options, specific to a DB2 for Linux, UNIX, and Windows solution, for operating system platforms and communications between warehousing components.	"Supported components"
Install prerequisite software before implementing your Tivoli Data Warehouse solution.	"Prerequisite installation" on page 541
Understand how to use the instructions for implementing your Tivoli Data Warehouse solution.	"Implementing a Tivoli Data Warehouse solution using DB2 for Linux, UNIX and Windows" on page 542
Complete the steps for implementing your Tivoli Data Warehouse solution using DB2 for the data warehouse.	<p>"Step 1: Create the Tivoli Data Warehouse database" on page 543</p> <p>"Step 2: Install and configure communications for the Warehouse Proxy Agent" on page 548</p> <p>"Step 3: Configure communications between the Tivoli Enterprise Portal Server and the data warehouse" on page 559</p> <p>"Step 4: Install and configure communications for the Summarization and Pruning Agent" on page 562</p> <p>"Step 5: Install and configure communications for the Tivoli Performance Analyzer" on page 563</p>

Supported components

This section illustrates the warehouse solution using DB2 for Linux, UNIX and Windows.

Figure 99 on page 540 presents the options for a Tivoli Data Warehouse solution using DB2 for Linux, UNIX and Windows for the warehouse database. The diagram summarizes the supported operating system platforms for the various warehousing components, the supported database products, and the connections between components. For more specific information about supported operating systems and database products, including product names and versions, see Chapter 3, "Hardware and software requirements," on page 75.

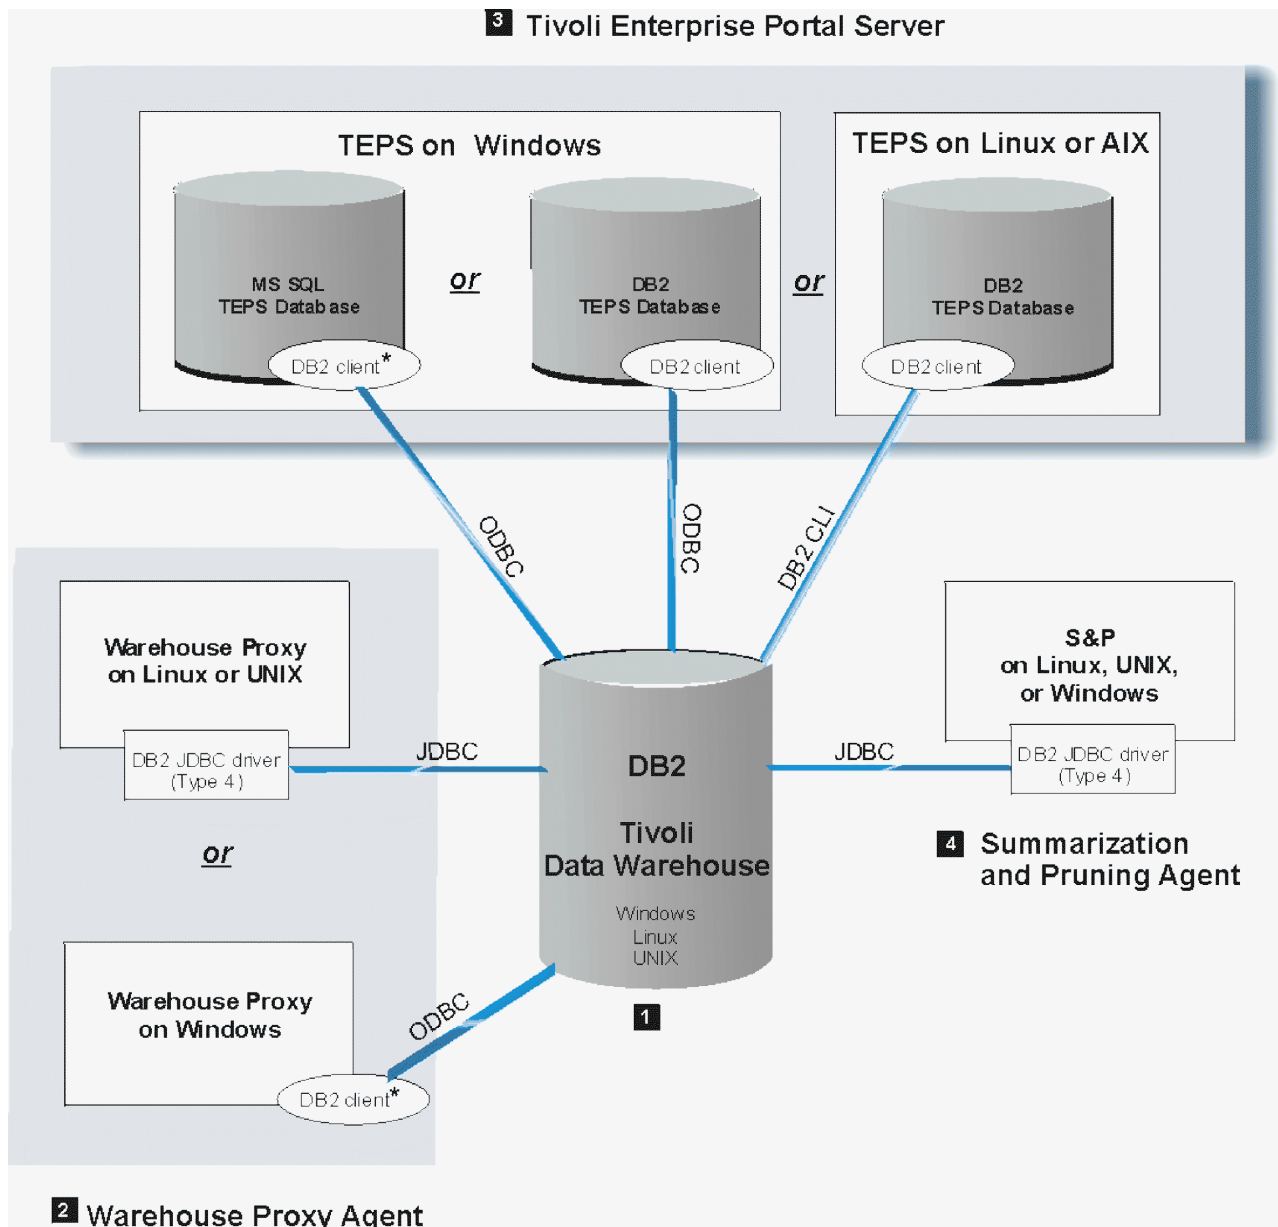


Figure 99. Tivoli Data Warehouse solution using DB2 for Linux, UNIX and Windows

Note: An asterisk (*) next to a database client indicates that you must manually install the client if it does not already exist.

In the following discussion, numbered product components correspond to the numbers on the diagram.

1 Tivoli Data Warehouse on DB2 for Linux, UNIX and Windows

A Tivoli Data Warehouse database on DB2 for Linux, UNIX and Windows can be installed on supported Windows, Linux, or any UNIX platform that is supported by DB2 for Linux, UNIX and Windows. Ensure that the DB2 TCP/IP listeners are active in order to accept connections from a DB2 client or JDBC driver.

2 Warehouse Proxy Agent

A Warehouse Proxy Agent on Linux or AIX communicates with the warehouse database through a JDBC connection. Install a Type 4 driver (DB2 for Linux, UNIX and Windows JDBC Universal Driver) on the computer where the Warehouse Proxy Agent is located.

A Warehouse Proxy Agent on Windows communicates with the warehouse database through an ODBC connection. The ODBC driver is included with the DB2 for Linux, UNIX and Windows client. If the Tivoli Data Warehouse is located on a remote computer, install a DB2 client on the local computer where the Warehouse Proxy Agent is located. Also, catalog the remote node and database on the local computer.

3 Tivoli Enterprise Portal Server

A Tivoli Enterprise Portal Server on Windows, Linux, or AIX can connect to a DB2 for Linux, UNIX and Windows data warehouse through a DB2 for Linux, UNIX and Windows client installed on the portal server. If the *portal server database* (designated as *TEPS database* in the diagram) uses DB2 for Linux, UNIX and Windows, the DB2 client already exists. *Manually* install a DB2 for Linux, UNIX and Windows client on the portal server only if the portal server database uses Microsoft SQL Server.

A portal server on Windows communicates with the warehouse database through an ODBC connection. The ODBC driver is included with the DB2 for Linux, UNIX and Windows client. Catalog the remote warehouse node and database on the portal server.

A portal server on Linux or AIX uses the DB2 for Linux, UNIX and Windows CLI interface, a proprietary connection, to communicate with the warehouse database. If the warehouse database is installed on a different computer from the portal server, catalog the remote database and remote node on the portal server.

4 Summarization and Pruning Agent

The Summarization and Pruning Agent communicates with the warehouse database through a JDBC connection from any supported operating system. Install a DB2 for Linux, UNIX and Windows Type 4 JDBC driver (*DB2 for Linux, UNIX and Windows JDBC Universal Driver*) on the computer where the Summarization and Pruning Agent is located.

Prerequisite installation

Before you implement your Tivoli Data Warehouse solution, complete one or more hub installations, *excluding the warehousing components*.

Include the following components in each hub installation:

- The hub Tivoli Enterprise Monitoring Server
- (Optional) One or more remote monitoring servers
- The Tivoli Enterprise Portal Server, including the prerequisite RDBMS for the portal server database (DB2 for Linux, UNIX and Windows or Microsoft SQL Server)
- An IBM DB2 for Linux, UNIX and Windows server on the computer where you will create the Tivoli Data Warehouse database. (The Tivoli Data Warehouse database can be shared in a multi-hub installation or dedicated to a single hub.)
- (Optional) A portal desktop client
- (Optional) Monitoring agents, and the application support for the monitoring agents

Note: The term *monitoring agent*, as used here, refers to agents that collect data directly from managed systems, not the Warehouse Proxy Agent or Summarization and Pruning Agent.

- (Optional) The Tivoli Performance Analyzer
- (Optional) Language packs for all languages other than English

See Table 82 for related information:

Table 82. Information topics related to installation of prerequisite software for a Tivoli Data Warehouse solution

Topic	Where to find information
Single and multiple hub installations	To understand the terminology related to single and multiple hub installations, see “Locating and sizing the hub Tivoli Enterprise Monitoring Server” on page 35.
Installation procedures for prerequisite components	The detailed instructions for installing the prerequisite components are described in Chapter 10, “Installing IBM Tivoli Monitoring,” on page 193. See your database documentations for instructions on how to install a supported database server.
Supported RDBMS versions	For specific information about the supported database platforms for the portal server database and the Tivoli Data Warehouse, see Chapter 3, “Hardware and software requirements,” on page 75.

Implementing a Tivoli Data Warehouse solution using DB2 for Linux, UNIX and Windows

The implementation instructions in this chapter are based on certain assumptions.

Use the instructions in the remainder of this chapter to implement a Tivoli Data Warehouse solution using DB2 for Linux, UNIX and Windows for the data warehouse.

Assumptions

The implementation instructions are based on the following assumptions:

- You will create the Tivoli Data Warehouse database on a different computer from the Tivoli Enterprise Portal Server.
- You will create a single Tivoli Data Warehouse database, to be used either within a single hub installation or to be shared in a multi-hub installation. If you have multiple independent hub installations, repeat the implementation steps for each hub installation. (See “Locating and sizing the hub Tivoli Enterprise Monitoring Server” on page 35 for information about hub installations.)
- No assumption is made about where you will install the Warehouse Proxy Agent and Summarization and Pruning Agent. Either of these agents may be installed on the same computer as the Tivoli Data Warehouse or on a different computer.

Solution steps

To implement your Tivoli Data Warehouse solution using DB2 for Linux, UNIX and Windows, complete the four major steps described in the remaining sections of this chapter, in the order listed.

About this task

Each major step consists of a series of installation and configuration tasks, listed and described in a table. Use the step tables as a road map for implementing your solution. The step tables describe the tasks at a high level, account for variations among configuration options (such as which operating system is used for a component), and reference the appropriate sections for detailed implementation procedures. To implement your solution successfully:

- Perform the tasks in the order listed in the table.
- Do not skip a table to the procedures that follow it.

Be aware that some of the implementation procedures referenced in a table are included in this chapter and some are documented elsewhere. In some cases, the task is described in the table, without referencing a separate procedure. Read and follow all instructions in the tables.

Procedure

1. Create the Tivoli Data Warehouse database.
2. Install and configure communications for the Warehouse Proxy Agent.
3. Configure communications between the Tivoli Enterprise Portal Server and the data warehouse.
4. Install and configure communications for the Summarization and Pruning Agent.

Step 1: Create the Tivoli Data Warehouse database

This section provides guidelines for creating a Tivoli Data Warehouse database using DB2 for Linux, UNIX and Windows.

Complete the tasks described in the following table to create a Tivoli Data Warehouse database using DB2 for Linux, UNIX and Windows and to make it accessible to clients.

Table 83. Tasks for creating the Tivoli Data Warehouse database

Task	Procedure
Create the Tivoli Data Warehouse database on one of the supported Windows, Linux, or UNIX operating systems. To comply with the assumptions described in the introduction to this chapter, create the warehouse database on a different computer from the Tivoli Enterprise Portal Server.	For guidance on planning the size and disk requirements for the warehouse database, see “Planning considerations for the Tivoli Data Warehouse” on page 510. For information about creating the warehouse database using DB2, see “Creating the warehouse database on DB2 for Linux, UNIX and Windows” on page 544.
Create an operating system (OS) user account (user name and password) with administrator authority on the computer where the warehouse database is located. The warehousing components (portal server, Warehouse Proxy Agents, and Summarization and Pruning Agent) will use this OS user account to access the database. Create only one user account for all warehousing components to use. This user is referred to in this chapter as the <i>warehouse user</i> .	“Creating a warehouse user on Windows” on page 544 “Creating a warehouse user on Linux or UNIX” on page 545
(Optional) Restrict the authority of the warehouse user. Initially, the warehouse user is created with OS administrative authority. Use the referenced procedure if you want to limit the authority of the warehouse user to just those privileges required to access and use the data warehouse.	“Limiting the authority of the warehouse user” on page 545
(Tivoli Data Warehouse on Linux or AIX only) Activate the DB2 for Linux, UNIX and Windows TCP/IP listeners on the DB2 server where the Tivoli Data Warehouse is installed. The TCP/IP listener processes on the DB2 for Linux, UNIX and Windows server must be active in order to accept connections from a DB2 for Linux, UNIX and Windows client or JDBC driver. The DB2 listeners are automatically activated on Windows systems. Perform the referenced procedure to activate the DB2 listeners if the warehouse database is located on a Linux or AIX system.	“Activating the DB2 listeners on a UNIX DB2 server” on page 547

Creating the warehouse database on DB2 for Linux, UNIX and Windows

This section provides guidelines for creating the Tivoli Data Warehouse database on DB2 for Linux, UNIX and Windows.

For specific instructions on how to create a DB2 database, see the DB2 for Linux, UNIX and Windows documentation or have a database administrator create the database for you.

When you create the warehouse database using DB2, follow these guidelines:

- Create the database with UTF-8 encoding.
- Create a name for the warehouse database, and an operating system (OS) user account (user name and password) that the warehousing components (portal server, Warehouse Proxy Agent, and Summarization and Pruning Agent) can use to access the data warehouse. In these instructions, this user account is referred to as the *warehouse user*.
- Consider using the default values shown in Table 84 for the warehouse name and warehouse user. The default values are used in the configuration procedures for connecting the warehousing components to the warehouse database. (For example, see the entries in Figure 101 on page 554.)

Table 84. Default values for Tivoli Data Warehouse parameters

Parameter	Default value
Tivoli Data Warehouse database name	WAREHOUS
User name	itmuser
User password	itmpswd1

- Give the warehouse user administrative authority to the database initially. After that, you can optionally limit the authority of the warehouse user to just the privileges required for interacting with the data warehouse. See the following sections for information about creating and limiting the authority of the warehouse user.
 - “Creating a warehouse user on Windows”
 - “Creating a warehouse user on Linux or UNIX” on page 545
 - “Limiting the authority of the warehouse user” on page 545
- For a Tivoli Data Warehouse on Linux or AIX, ensure that the DB2 for Linux, UNIX and Windows server is configured for TCP/IP communications. See “Activating the DB2 listeners on a UNIX DB2 server” on page 547.

Creating a warehouse user on Windows

Complete the steps in this section on the computer where the warehouse database is installed to create a Windows OS user with Administrator authority.

Procedure

1. Right-click the **My Computers** icon on the Windows desktop and click **Manage**.
2. In the navigation pane of the Computer Management window, expand **Local Users and Groups** by clicking on the plus sign (+).
3. Right-click the **Users** folder and click **New User**.
4. Type a user name and password in the **User Name** and **Password** fields. Confirm the password by typing it again in the **Confirm password** field.
5. Clear **User must change password at next logon**.
6. Click **Close**.
7. Click the **Groups** folder.

8. Double-click **Administrators** in the right pane of the window.
9. Click **Add** in the Administrator Properties window.
10. Locate the new user you created and select it.
11. Click **Add**.
12. Click **OK** and then **OK** again to close the Administrator Properties window.
13. Close the Computer Management window.

Creating a warehouse user on Linux or UNIX

Complete the procedure in this section on the computer where the warehouse database is installed to create a Linux or UNIX OS user with administrative authority to the warehouse.

- To create the user, follow the instructions in the documentation for the specific Linux or UNIX product and version that is installed on the computer where the warehouse database is located.
- To give this user administrative authority to the data warehouse, add the user to the DB2 for Linux, UNIX and Windows SYSADM group. Run the following command to find the name of the SYSADM group:

```
db2 get dbm cfg | grep SYSADM
```

For example:

```
db2 get dbm cfg | grep SYSADM
SYSADM group name          (SYSADM_GROUP) = DB2GRP1
```

In this example, the name of the DB2 for Linux, UNIX and Windows SYSADM group is DB2GRP1. If you created an OS user named ITMUSER, add ITMUSER to DB2GRP1.

Limiting the authority of the warehouse user

If you do not want the warehouse user to have broad administrative authority, you can limit the authority of the warehouse user to just those privileges required for accessing and using the data warehouse. These more limited privileges include the authority to create and update tables, to insert or delete information from the tables, to create indexes for the tables, and to grant public authority to the tables.

About this task

The Tivoli Data Warehouse requires one bufferpool and three tablespaces to begin its operation. The bufferpool and tablespaces are created by the warehouse user before the Warehouse Proxy Agent starts, provided the warehouse user has administrative authority to the database. A warehouse user with limited authority cannot create the required bufferpool and tablespaces. Therefore, this procedure to limit the authority of the warehouse user includes steps to create the bufferpool and tablespaces in advance. Be sure to perform this procedure before the Warehouse Proxy Agent starts.

Use the script shown in the next table to create the required bufferpool and tablespaces. Create the script on a computer from which you can connect to the Tivoli Data Warehouse database server.

```
-- CREATE a Bufferpool of page size 8K
CREATE BUFFERPOOL ITMBUF8K IMMEDIATE SIZE 2501 PAGESIZE 8K;

-- CREATE a Regular Tablespace using the 8K Bufferpool
CREATE REGULAR TABLESPACE ITMREG8K PAGESIZE 8K
MANAGED BY SYSTEM
USING ('itmreg8k')2 BUFFERPOOL ITMBUF8k;

-- CREATE a System tablespace using the 8K Bufferpool
CREATE SYSTEM TEMPORARY TABLESPACE ITMSYS8K PAGESIZE 8K
MANAGED BY SYSTEM
USING ('itmsys8k')2 BUFFERPOOL ITMBUF8k;

-- CREATE a User tablespace using the 8K Bufferpool
CREATE USER TEMPORARY TABLESPACE ITMUSER8K PAGESIZE 8K
MANAGED BY SYSTEM
USING ('itmuser8k')2 BUFFERPOOL ITMBUF8k;
```

Note:

1. SIZE is the number of 8K pages to allocate for the bufferpool. If there is sufficient memory, performance can be improved by increasing this number.
2. A fully qualified path can be specified here. As shown this will be created in a default directory.
3. IBM Tivoli Monitoring creates SMS tablespaces, which cannot be extended across multiple drives. This can be customized by the database administrator.




To limit the authority of the warehouse user, complete the following steps:

Procedure

1. Connect to the data warehouse with db2admin privileges:

```
db2 connect to warehouse user db2admin using password
```

where *warehouse* is the name of the Warehouse database, *db2admin* is the DB2 for Linux, UNIX and Windows administrator ID, and *password* is the password of the *db2admin* user ID. The user ID must be a DB2 user with SYSADM authority

2. Change to the directory where the KHD_DB2_crt_BP_TBSP.sql script is located.
3. Run the script to create the required bufferpool and tablespaces:
db2 -stvf KHD_DB2_crt_BP_TBSP.sql
4. Remove administrative privileges from the warehouse user (OS user) that you created when you created the warehouse database:
 -  Remove the warehouse user from the Administrator group.
 -   Remove the warehouse user from the SYSADM group to which it was assigned (for example, DB2GRP1). (See “Creating a warehouse user on Linux or UNIX” on page 545.)
5. Grant these database authorities to the warehouse user:
 - CONNECT
 - CREATETAB

Note: The user might already implicitly have CONNECT and CREATEAB authority through the PUBLIC role.

- USE OF TABLESPACE

CONNECT authority grants the user access to the database. CREATETAB authority allows the user to create tables. The authority to drop tables, alter tables, create and drop indexes for the tables, insert, delete, or update data in the tables, are all implicitly granted. USE OF TABLESPACE grants the user authority to use particular tablespaces, in this case ITMREG8K.

To grant these authorities, you can use either the DB2 Control Center (Database Authorities window) or the command-line interface. If you use the command-line interface, run commands similar to the following. In this example, the name of the warehouse user is `itmuser`.

```
db2 "GRANT CONNECT ON DATABASE TO USER itmuser"
db2 "GRANT CREATETAB ON DATABASE TO USER itmuser"
db2 "GRANT USE OF TABLESPACE ITMREG8K TO USER itmuser"
```

What to do next

Perform the additional steps if your security policy prohibits use of the `CREATTAB` authority:

1. Use the schema tool to generate the DDL that create the database objects.

Note: You should create the historical collections that you want first and then configure the Summarization and Pruning Agent so that you can use the schema tool's configured mode.

2. Execute the generated scripts as described in Chapter 20, "Schema Publication Tool," on page 525. Run the scripts as the Tivoli Data Warehouse user so that the Tivoli Data Warehouse user has sufficient privileges on the tables.
3. Revoke the `CREATETAB` privilege from the Tivoli Data Warehouse user and from the `PUBLIC` role.

Setting database and instance configuration values

The default values for many of the configuration values for the DB2 for Linux, UNIX and Windows database and the DB2 for Linux, UNIX and Windows instance may not be acceptable for a production environment.

Value	Purpose
<code>LOGPRIMARY</code>	Number of primary transaction logs. In a production environment, this value should be larger than the default value.
<code>LOGSECOND</code>	Number of secondary log files. In a production environment, this value should be larger than the default value.
<code>NEWLOGPATH</code>	Location of transaction logs. To improve performance, the logs should be placed on a separate physical disk from the tables and indexes.
<code>LOGFILSIZ</code>	Size of transaction log files. In a production environment, this value should be larger than the default value.
<code>LOCKTIMEOUT</code>	Defaults to infinite wait, which can cause Warehouse Proxy and Summarization and Pruning processing to appear to be locked up. This should be set to a reasonable value such as 120 seconds.

Activating the DB2 listeners on a UNIX DB2 server

The TCP/IP listener processes on the DB2 for Linux, UNIX and Windows server where the Tivoli Data Warehouse database is installed must be active in order to accept connections from a DB2 for Linux, UNIX and Windows client or a JDBC Type 4 driver (*DB2 for Linux, UNIX and Windows JDBC Universal Driver*).

On a Windows system, the DB2 listeners are automatically activated. Run the following commands on a UNIX system where the Tivoli Data Warehouse database is installed to activate the DB2 for Linux, UNIX and Windows listeners:

```
db2set -i instance_name DB2COMM=tcPIP
db2 update dbm cfg using SVCENAME port_number
db2stop
db2start
```

where *instance_name* is the name of the instance in which you created the warehouse database and *port_number* is the listening port for the instance. (The port number is specified in the file /etc/services.) For example:

```
db2set -i db2inst1 DB2COMM=tcPIP
db2 update dbm cfg using SVCENAME 60000
db2stop
db2start
```

Step 2: Install and configure communications for the Warehouse Proxy Agent

You can install one or more Warehouse Proxy Agents to collect and send historical data to the Tivoli Data Warehouse database.

Complete the tasks described in the following table, in the order listed, to install and configure each Warehouse Proxy Agent.

Table 85. Tasks for installing and configuring communications for the Warehouse Proxy Agent

Task	Procedure
<p>Install one or more Warehouse Proxy Agents. If you want to install a Summarization and Pruning Agent on the same computer as one of the Warehouse Proxy Agents, use the referenced procedures to install both agents at the same time.</p> <p>If you are installing more than one Warehouse Proxy Agent, each agent must be installed on a separate computer.</p> <p>The installation procedure for Windows includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring server. On Linux or AIX, this step is performed in a separate configuration procedure (<i>Configuring the monitoring agent</i>) and an X11 GUI is required to configure the agent. Alternatively, you can run the following command to utilize an X terminal emulation program (such as Cygwin) that is running on another computer:</p> <pre>export DISPLAY=my_windows_pc_IP_addr:0.0</pre> <p>where <i>my_windows_pc_IP_addr</i> is the IP address of a computer that is running an X terminal emulation program. See the information at right. Be sure to perform all referenced installation and configuration procedures.</p> <p>Note for sites setting up autonomous operation: The installation procedure includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring Server. On Windows operating systems, if you want to run the Warehouse Proxy Agent without a connection to the hub, accept the defaults for the connection information, but specify a nonvalid name for the monitoring server. On UNIX and Linux operating systems, check No TEMS on the TEMS Connection tab of the configuration window.</p>	<p>To install a Warehouse Proxy Agent on Windows, complete the procedure “Windows: Installing a monitoring agent” on page 262.</p> <p>To install a Warehouse Proxy Agent on Linux or AIX, complete the procedure “Linux or UNIX: Installing a monitoring agent” on page 267, including the following subsections:</p> <ul style="list-style-type: none"> • <i>Installing the monitoring agent</i> • <i>Configuring the monitoring agent</i> • <i>Changing the file permissions for agents</i> (if you used a non-root user to install the Warehouse Proxy) <p><i>Do not complete the procedure for starting the agent.</i></p>

Table 85. Tasks for installing and configuring communications for the Warehouse Proxy Agent (continued)

Task	Procedure
<p>(Warehouse Proxy Agent on Windows only)</p> <ul style="list-style-type: none"> Install a DB2 for Linux, UNIX and Windows client on the computer where the Warehouse Proxy Agent is installed if <i>both</i> of the following statements are true: <ul style="list-style-type: none"> The Warehouse Proxy is installed on Windows, and The Warehouse Proxy needs to connect to a remote data warehouse. Catalog the remote data warehouse on the Windows computer where you installed the DB2 for Linux, UNIX and Windows client. You must perform this step before configuring an ODBC data source. (See the next row.) Set the following system variable on the computer where the Warehouse Proxy Agent is installed. Restart the computer after setting the variable. DB2CODEPAGE=1208 Set the environment variable whether or not the warehouse database is local or remote. 	<p>See the DB2 for Linux, UNIX and Windows documentation for instructions on how to install a DB2 for Linux, UNIX and Windows client.</p> <p>To catalog a remote data warehouse, see “Cataloging a remote data warehouse” on page 550.</p>
<p>(Warehouse Proxy Agent on Windows only)</p> <p>On the computer where the Warehouse Proxy Agent is installed, configure an ODBC data source for the data warehouse.</p> <p>Perform this procedure whether or not the Warehouse Proxy Agent and the warehouse database are installed on the same computer.</p>	<p>“Configuring an ODBC data source for a DB2 data warehouse” on page 551</p>
<p>(Warehouse Proxy Agent on Linux or AIX only)</p> <p>If the data warehouse is located on a remote computer, copy the <i>DB2 for Linux, UNIX and Windows JDBC Universal Driver</i> (Type 4 driver) JAR files, included with the DB2 for Linux, UNIX and Windows product installation, to the local computer where the Warehouse Proxy Agent is installed. You can copy the files to any directory on the local computer.</p>	<p>The Type 4 driver file names and locations are as follows:</p> <pre>db2installdir/java/db2jcc.jar db2installdir/java/db2jcc_license_cu.jar</pre> <p>where <i>db2installdir</i> is the directory where DB2 for Linux, UNIX and Windows was installed. The default DB2 for Linux, UNIX and Windows Version 9 installation directory is as follows:</p> <ul style="list-style-type: none"> On AIX: /usr/opt/db2_09_01 On Linux: /opt/IBM/db2/V9.1
<p>Configure the Warehouse Proxy Agent to connect to the data warehouse.</p> <p>Perform this procedure whether or not the Warehouse Proxy Agent and the warehouse database are installed on the same computer.</p>	<p>For a Warehouse Proxy Agent on Windows, see “Configuring a Warehouse Proxy Agent on Windows (ODBC connection)” on page 552.</p> <p>For a Warehouse Proxy Agent on Linux or AIX, see “Configuring a Warehouse Proxy Agent on Linux or UNIX (JDBC connection)” on page 555.</p>
<p>If you are installing more than one Warehouse Proxy Agent within the same hub monitoring server installation, associate each Warehouse Proxy Agent with a subset of monitoring servers (hub or remote) within the installation. Each Warehouse Proxy Agent receives data from the monitoring agents that report to the monitoring servers on the list. Use the environment variable KHD_WAREHOUSE_TEMS_LIST to specify a list of monitoring servers to associate with a Warehouse Proxy Agent.</p>	<p>For instructions about installing and configuring multiple Warehouse Proxy Agents within a single hub monitoring server installation, see “Installing and configuring multiple Warehouse Proxy Agents” on page 656.</p>

Table 85. Tasks for installing and configuring communications for the Warehouse Proxy Agent (continued)

Task	Procedure
(Optional) Customize the configuration of the Warehouse Proxy Agent for tuning performance.	"Tuning the performance of the Warehouse Proxy" on page 666
Start the Warehouse Proxy Agent.	"Starting the Warehouse Proxy Agent" on page 558

Cataloging a remote data warehouse

Perform this procedure on a computer where a DB2 for Linux, UNIX and Windows client is installed to enable communication between the client and a *remote* DB2 for Linux, UNIX and Windows server where the data warehouse is installed.

About this task

For example, use this procedure to set up communication to a remote DB2 data warehouse server from:

- The DB2 for Linux, UNIX and Windows client on the computer where the Tivoli Enterprise Portal Server is installed (on any platform)
- The DB2 client on a Windows computer where a Warehouse Proxy Agent installed

Do *not* perform this procedure on the computer where the data warehouse (DB2 server) is installed or on a computer where there is no DB2 client (for example, on a computer where a Type 4 DB2 for Linux, UNIX and Windows JDBC driver is used to communicate with the remote DB2 server).

Complete the following steps on the computer where the DB2 for Linux, UNIX and Windows client is installed (the *local* computer):

Procedure

1. Catalog the remote TCP/IP node where the warehouse database is installed:

```
db2 catalog tcpip node node_name remote host_name server port
db2 terminate
```

where the indicated variables identify the location and port of the remote DB2 for Linux, UNIX and Windows server. For *host_name*, specify the host name or IP address. The default port for a DB2 server is 60000. For example:

```
db2 catalog tcpip node amsc2424 remote 8.53.36.240 server 60000
db2 terminate
```

2. Catalog the remote Tivoli Data Warehouse database:

```
db2 catalog db db_name as db_alias at node node_name
db2 terminate
```

where

db_name is the name of the remote warehouse database.

db_alias is the nickname or alias used to identify the remote warehouse database on the local computer. The local alias for the warehouse database must match the name that you specify in the configuration procedure for the portal server, Warehouse Proxy Agent, or Summarization and Pruning Agent.

node_name is the name of the node where the warehouse database is located.

Example:

```
db2 catalog db WAREHOUS as WAREHOUS at node amsc2424
db2 terminate
```

3. Test the connection to the remote warehouse database:

```
db2 connect to db_alias user user_name using user_password
```

where:

db_alias

Is the nickname or alias used to identify the remote warehouse database on the local computer.

user_name

Is the user ID that the local DB2 client uses to access the warehouse database.

user_password

Is the password for that *user_name*.

These values must match the values that you specify in the configuration procedures for the portal server, Warehouse Proxy Agent, or Summarization and Pruning Agent.

Example:

```
db2 connect to WAREHOUS user itmuser using itmpswd1
```

Configuring an ODBC data source for a DB2 data warehouse

A DB2 for Linux, UNIX and Windows client on Windows requires an ODBC connection to the data warehouse. For the Warehouse Proxy Agent, you must configure the ODBC connection manually.

Before you begin

- If the warehouse database is remote from the Warehouse Proxy Agent, catalog the remote database before you configure the ODBC data source. See “Cataloging a remote data warehouse” on page 550. The ODBC connection does not work if the remote database has not been cataloged prior to performing this procedure.
- This procedure uses default values for the data source name, warehouse alias, and warehouse user ID. (Default values are used in configuration procedures for warehousing components.) Substitute different values if you do not want to use the default values.

About this task

Complete the following procedure to set up an ODBC connection for a Warehouse Proxy Agent on Windows to a local or remote Tivoli Data Warehouse.

Procedure

1. On the computer where the Warehouse Proxy Agent is installed, open the Control Panel.
2. Click **Administrative Tools** → **Data Sources (ODBC)**
3. Click **Add** in the **System DSN** tab in the ODBC Data Source Administrator window.
4. Select **IBM DB2 ODBC DRIVER** from the list.
5. Click **Finish**.
6. In the ODBC DB2 Driver - Add window, perform the following steps:
 - a. Enter ITM Warehouse in **Data source name**.
 - b. Enter Warehous in **Database Alias**.

If the Tivoli Data Warehouse is located on a remote computer, ensure that the database alias matches the alias that you used when cataloging the remote data warehouse. See “Cataloging a remote data warehouse” on page 550.

If local, ensure that the database alias matches the name used for the warehouse database.
 - c. Click **OK**.
7. Test the ODBC database connection before continuing:
 - a. In the ODBC Data Source Administrator window, select **ITM Warehouse**.
 - b. Click **Configure**.

- c. In the CLI/ODBC Settings - ITM Warehouse window, you see the data source name, **ITM Warehouse**.
- d. Enter ITMUser for the **User ID**.
- e. Type a password for the user in the **Password** field. The default password is itmpswd1.
- f. Click **Connect**.
- g. A **Connection test successful** message is displayed.
- h. Click **OK**.
- i. Click **OK** to close the window.

Configuring a Warehouse Proxy Agent on Windows (ODBC connection)

Use the procedure in this section to configure a Warehouse Proxy Agent on Windows to connect to a DB2 for Linux, UNIX, and Windows data warehouse.

Procedure

1. Log on to the Windows system where the Warehouse Proxy Agent is installed and begin the configuration:
 - a. Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Monitoring Services**.
The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Warehouse Proxy** and click **Configure Using Defaults**.
Click **Reconfigure** if the Warehouse Proxy is installed on the same computer as the portal server.
 - c. Click **OK** on the message regarding connection to a hub monitoring server.
2. The next two windows (entitled Warehouse Proxy: Agent Advanced Configuration) contain the settings for the connection between the Warehouse Proxy Agent and the hub monitoring server. These settings were specified when the Warehouse Proxy Agent was installed. Click **OK** on each window to accept the settings.
3. Select **DB2** from the list of selectable databases (see Figure 100 on page 553), and click **Next**.




Figure 100. Warehouse Proxy Database Selection screen

The following configuration window is displayed.

Figure 101. Configure DB2 Data Source for Warehouse Proxy window

4. Click **OK** to accept all default information on this window, or change one or more default values and then click **OK**. The fields on this window are described in Table 86.

Note: The values for the data source name, database name, and database user ID and password must match the values that you used when configuring an ODBC connection for the Warehouse Proxy Agent. (See “Configuring an ODBC data source for a DB2 data warehouse” on page 551.)

Table 86. Configuration information for the Tivoli Data Warehouse database on DB2 for Linux, UNIX and Windows

Field	Default value	Description
ODBC DSN	ITM Warehouse	The name of the data source.
Username	ITMUser	The name of the user that the Warehouse Proxy Agent will use to access the Tivoli Data Warehouse database.
Password	itmpswd1	The password that the Warehouse Proxy Agent will use to access the Tivoli Data Warehouse database. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Confirm Password	itmpswd1	Confirm the password by entering it again.
Test Connection		Test the connection to the Tivoli Data Warehouse Database based on the completed fields above: ODBC DSN, Username, and Password. Note: Test Connection is not available if configuring remotely from the Tivoli Enterprise Portal.

Table 86. Configuration information for the Tivoli Data Warehouse database on DB2 for Linux, UNIX and Windows (continued)

Field	Default value	Description
Warehouse TEMS List		Environment variable containing a space delimited list of TEMS names, which are given during the configuration of HTEMS or RTEMS. A TEMS name in this field indicates that all the agents connected to this TEMS will have their historical data sent to this Warehouse Proxy Agent. This variable is used when the ITM environment contains multiple Warehouse Proxy Agents and the workload has to be balanced using specific Warehouse Proxy Agents.
Use Batch		Batch inserts can greatly increase the data insertion rate of the Warehouse Proxy Agent. This is especially true if the proxy and the warehouse are located on different hosts. Batch inserts are supported for ODBC warehouse connections. Using batch inserts is recommended in all configurations, but they will place increased load on the data warehouse.
Database Compression		If the database compression mode is supported by the database, the Warehouse Proxy Agent will create all tables and indexes in the Tivoli Data Warehouse with compression enabled. This option reduces the storage costs of the Tivoli Data Warehouse.
Warehouse Compression for Z/OS Sources		Select this option for the Warehouse Proxy server to allow clients installed on Z/OS machines to send compressed data.
Warehouse Compression for Distributed Sources		Select this option for the Warehouse Proxy server to allow clients installed on distributed machines (Linux/UNIX, Windows) to send compressed data.

5. Click OK.

Configuring a Warehouse Proxy Agent on Linux or UNIX (JDBC connection)

Complete the procedure in this section to configure a Warehouse Proxy Agent on Linux or UNIX to connect to a DB2 for Linux, UNIX and Windows Tivoli Data Warehouse on any operating system.

Procedure

1. Log on to the computer where the Warehouse Proxy Agent is installed and begin the configuration.
 - a. Change to the *install_dir/bin* directory and run the following command:


```
./itmcmd manage [-h install_dir]
```

where *install_dir* is the installation directory for IBM Tivoli Monitoring. The default installation directory is /opt/IBM/ITM.

The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Warehouse Proxy** and click **Configure**.
The Configure Warehouse Proxy window is displayed.




Figure 102. Configure Warehouse Proxy window (Database Type)

2. Select **DB2** from the list of selectable databases, and click **Next**. The following configuration window is displayed.

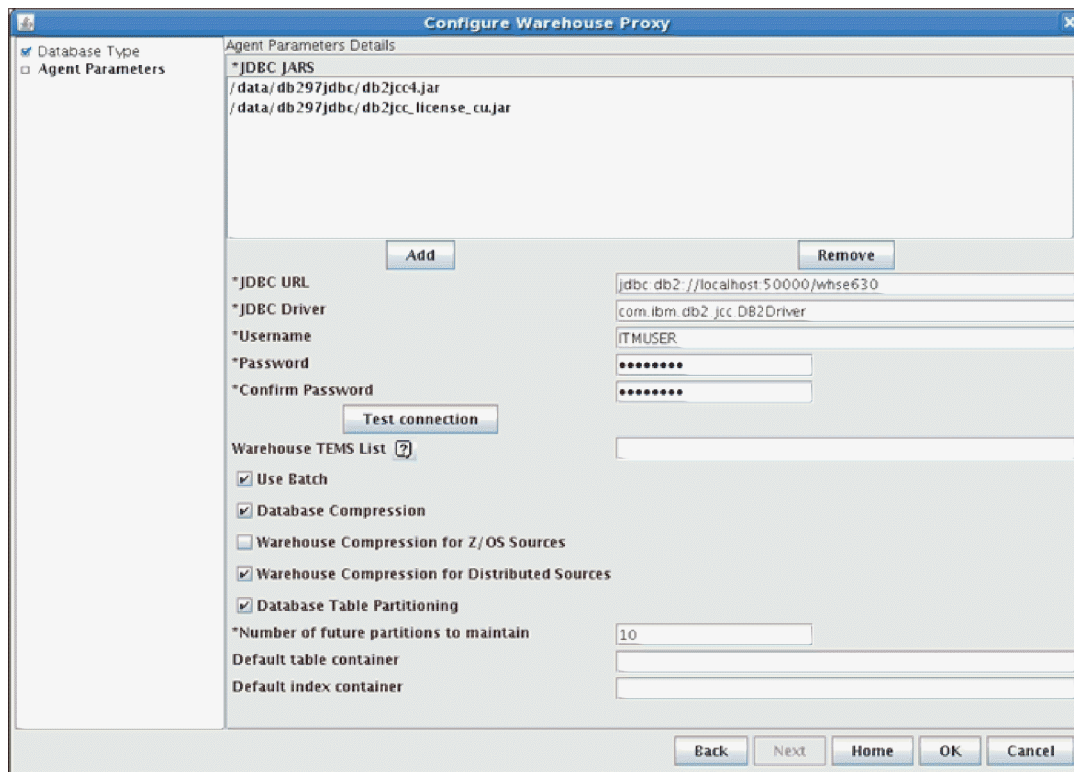


Figure 103. Configure Warehouse Proxy window (Agent Parameters)

3. Review the settings for the connection between the Warehouse Proxy Agent and the hub monitoring server. Correct the settings if necessary.

The Warehouse Proxy Agent must use the same protocols used by the application agents and by the hub monitoring. If the proxy agent does not have the same protocol as the hub monitoring server, it cannot register with the hub. If the proxy does not have the same protocol as the application agents, then the application agents cannot communicate with the proxy when they try to create a route to it.

4. Add the names and directory locations of the JDBC JARS files to the **JDBC JARS** list box:
 - a. Click **Add** to display the file browser window. Navigate to the location of the JARS files on this computer and select the following driver files:

db2jcc.jar
db2jcc_license_cu.jar

- b. Click **OK** to close the browser window and add the JDBC JARS files to the list.

If you need to delete an entry from the list, select the entry and click **Remove**.

5. Change the default value displayed in the **JDBC URL** field if it is not correct. The default Tivoli Data Warehouse URL for IBM DB2 for Linux, UNIX and Windows is as follows:

jdbc:db2://localhost:50000/WAREHOUS

- If the Tivoli Data Warehouse is installed on a remote computer, specify the host name of the remote computer instead of localhost.
- Change the port number if it is different.
- If the name of the Tivoli Data Warehouse database is not WAREHOUS, replace WAREHOUS with the actual name. (See “Creating the warehouse database on DB2 for Linux, UNIX and Windows” on page 544.)

Note: When specifying the DB2 for Linux, UNIX and Windows database name on Linux and UNIX, case is ignored. In other words, it makes no difference whether you provide the database name in lowercase or uppercase.

6. Verify the JDBC driver name, which is displayed in the **JDBC Driver** field. (Note that the **JDBC Driver** field displays the *driver name*, in contrast to the *JDBC JARS* that are listed in the **JDBC JARS** field.)

The default DB2 JDBC Driver is:

`com.ibm.db2.jcc.DB2Driver`

7. If necessary, change the entries in the **Username** and **Password** fields to match the user name and password that were created for the Tivoli Data Warehouse. (See “Creating the warehouse database on DB2 for Linux, UNIX and Windows” on page 544.) The default user name is `itmuser` and the default password is `itmpswd1`.

8. Click **Test connection** to ensure you can communicate with the Tivoli Data Warehouse database.

Note: Test Connection is not available if configuring remotely from the Tivoli Enterprise Portal.

9. A Tivoli Enterprise Monitoring Server name in the **Warehouse TEMS List** field indicates that all the agents connected to this Tivoli Enterprise Monitoring Server will have its historical data sent to this Warehouse Proxy Agent. This variable is used when the IBM Tivoli Monitoring environment contains multiple Warehouse Proxy Agents and the workload has to be balanced using specific Warehouse Proxy Agents.
10. Select the **Use Batch** check box if you want the Warehouse Proxy Agent to submit multiple execute statements to the Tivoli Data Warehouse database for processing as a batch.
In some situations, such as crossing a network, sending multiple statements as a unit is more efficient than sending each statement separately. Batch processing is one of the features provided with the JDBC 2.0 API.
11. Select the **Database Compression** check box for the Warehouse Proxy Agent to create all tables and indexes in the Tivoli Data Warehouse with compression enabled, if the database compression mode is supported by the database. This option reduces the storage costs of the Tivoli Data Warehouse.
12. Select the **Warehouse Compression for Z/OS Sources** check box for the Warehouse Proxy server to allow clients installed on Z/OS machines to send compressed data.
13. Click **OK** to save your settings and close the window.

Starting the Warehouse Proxy Agent

This section contains instructions for starting the Warehouse Proxy Agent.

- To start the Warehouse Proxy Agent from the Manage Tivoli Enterprise Monitoring Services window, right-click **Warehouse Proxy** and select **Start**.
- (*Linux or AIX only*) To start the Warehouse Proxy Agent from the command-line, run the following command from the bin directory of the IBM Tivoli Monitoring installation directory. The default installation directory is `/opt/IBM/ITM`.
`./itmcmd agent start hd`

where `hd` is the product code for the Warehouse Proxy Agent.

Step 3: Configure communications between the Tivoli Enterprise Portal Server and the data warehouse

Complete the tasks described in this section, in the order listed, to configure communications between the portal server and the data warehouse.

Table 87. Tasks for configuring communications between the portal server and a DB2 for Linux, UNIX and Windows data warehouse

Task	Procedure
<i>(Portal server on Windows only)</i> If the <i>portal server database</i> was created on Microsoft SQL Server, install a DB2 for Linux, UNIX and Windows database client on the portal server. If the portal server database was created on DB2 for Linux, UNIX and Windows, the DB2 client already exists on the portal server.	See the DB2 for Linux, UNIX and Windows documentation for instructions on how to install a DB2 for Linux, UNIX and Windows client.
On the computer where the portal server is installed, catalog the remote data warehouse. You must perform this step before configuring the portal server to connect to the data warehouse. (See the next row.) Cataloging the remote data warehouse enables communications between the DB2 for Linux, UNIX and Windows client on the portal server and the remote DB2 for Linux, UNIX and Windows server where the data warehouse is located. Complete this task regardless of which platforms are used by the portal server or the data warehouse.	"Cataloging a remote data warehouse" on page 550
Configure the portal server to connect to the data warehouse. The configuration procedure on Windows automatically configures an ODBC connection to the data warehouse.	For a portal server on Windows, see "Configuring a Windows portal server (ODBC connection)." For a portal server on Linux or AIX, see "Configuring a Linux or AIX portal server (DB2 for Linux, UNIX and Windows CLI connection)" on page 561.
Restart the portal server.	"Starting the portal server" on page 562
Test the connection between the portal server and the Tivoli Data Warehouse by creating a customized query in the Tivoli Enterprise Portal.	"Testing the connection between the portal server and the Tivoli Data Warehouse" on page 663

Configuring a Windows portal server (ODBC connection)

The procedure described in this section uses the Manage Tivoli Enterprise Monitoring Services window to configure an ODBC connection between a Windows portal server and the data warehouse. You do not need to configure the ODBC data source through the Control Panel in Windows.

Before you begin

Catalog the remote warehouse database before you configure the Windows portal server to connect to the database. See "Cataloging a remote data warehouse" on page 550. The ODBC connection does not work if the remote database has not been cataloged prior to performing this procedure.

About this task

Complete the following procedure to configure a portal server on Windows to connect to a DB2 for Linux, UNIX and Windows data warehouse:

Procedure

- Log on to the Windows system where the portal server is installed and begin the configuration:
 - Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Monitoring Services**. The **Manage Tivoli Enterprise Monitoring Services** window is displayed.
 - Right-click **Tivoli Enterprise Portal Server** and click **Reconfigure**.
- The next two windows (entitled **TEP Server Configuration**) contain the settings for the connection between the portal server and the hub monitoring server. These settings were specified when the portal server was installed. Click **OK** on each window to accept the settings.
- Click **Yes** on the message asking if you want to reconfigure the warehouse information for the Tivoli Enterprise Portal Server.
- Select **DB2** from the list of databases and click **OK**.
The following configuration window is displayed.

Figure 104. Configure DB2 Data Source for Warehouse window

- Click **OK** to accept all default information on this window, or change one or more default values and then click **OK**. The fields on this window are described in the following table.

Table 88. Configuration information for the Tivoli Data Warehouse database on DB2 for Linux, UNIX and Windows

Field	Default value	Description
Data Source Name	ITM Warehouse	The name of the data source.
Database User ID	ITMUser	The name of the Windows OS user that the portal server will use to access the Tivoli Data Warehouse database.
Database Password	itmpswd1	The password for the Windows user. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Reenter Password	itmpswd1	Confirm the password by entering it again.

6. Click OK.

Configuring a Linux or AIX portal server (DB2 for Linux, UNIX and Windows CLI connection)

Use this procedure to configure a portal server on Linux or AIX to connect to a DB2 for Linux, UNIX and Windows Tivoli Data Warehouse on any operating system.

Procedure

1. Log on to the computer where the Tivoli Enterprise Portal Server is installed and begin the configuration.
 - a. Change to the *install_dir/bin* directory and run the following command:


```
./itmcmd manage [-h install_dir]
```

where *install_dir* is the installation directory for IBM Tivoli Monitoring. The default installation directory is /opt/IBM/ITM.

The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Tivoli Enterprise Portal Server** and click **Configure**.
The Configure Tivoli Enterprise Portal Server window is displayed.
2. On the TEMS Connection tab, review the settings for the connection between the portal server and the hub monitoring server. These settings were specified when the portal server was installed.
3. Click the **Agent Parameters** tab.
4. Select the **DB2** radio button.

The fields for configuring the connection to a DB2 for Linux, UNIX and Windows data warehouse are displayed at the bottom of the window.

Figure 105. Configuring the connection to a DB2 for Linux, UNIX and Windows data warehouse

5. Enter information in the fields described in the following table:

Table 89. Configuration information for the Tivoli Data Warehouse database on DB2 for Linux, UNIX and Windows

Field	Default value	Description
Warehouse database name	WAREHOUS	The name of the Tivoli Data Warehouse database.
Warehouse database user ID	itmuser	The login name of the database user that the portal server will use to access the Tivoli Data Warehouse database.
Warehouse user password	itmpswd1	The password for the database login user. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Re-type warehouse user password	itmpswd1	Confirm the password by entering it again.

6. Click **Save** to save your settings and close the window.

Starting the portal server

This section contains instructions for starting the portal server.

Use the following steps to start the portal server:

- To start the portal server from the Manage Tivoli Enterprise Monitoring Services window, right-click **Tivoli Enterprise Portal Server** and select **Start**.
- (*Linux or AIX only*) To start the portal server from the command-line, run the following command from the bin directory of the IBM Tivoli Monitoring installation directory. The default installation directory is /opt/IBM/ITM.
./itmcmd agent start cq

where cq is the product code for the portal server.

Step 4: Install and configure communications for the Summarization and Pruning Agent

Complete the tasks described in the following table, in the order listed, to install and configure the Summarization and Pruning Agent.

Table 90. Tasks for installing and configuring communications for the Summarization and Pruning Agent

Task	Procedure
<p>Install the Summarization and Pruning Agent if you have not already installed it. For best performance, install the Summarization and Pruning Agent on the same computer as the data warehouse.</p> <p>The installation procedure for Windows includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring server. On Linux or AIX, this step is performed in a separate configuration procedure (<i>Configuring the monitoring agent</i>). See the information at right. Be sure to perform all referenced installation and configuration procedures.</p> <p>Note: The Summarization and Pruning Agent is not automatically started after installation. Do not complete any step or procedure for starting the agent at this point.</p>	<p>To install a Summarization and Pruning Agent on Windows, complete the procedure “Windows: Installing a monitoring agent” on page 262.</p> <p>To install a Summarization and Pruning Agent on Linux or UNIX, complete the procedure “Linux or UNIX: Installing a monitoring agent” on page 267, including the following subsections:</p> <ul style="list-style-type: none"> • <i>Installing the monitoring agent</i> • <i>Configuring the monitoring agent</i> • <i>Changing the file permissions for agents</i> (if you used a non-root user to install the Warehouse Proxy) <p><i>Do not complete the procedure for starting the agent.</i></p>

Table 90. Tasks for installing and configuring communications for the Summarization and Pruning Agent (continued)

Task	Procedure
If the data warehouse is located on a remote computer, copy the <i>DB2 for Linux, UNIX and Windows JDBC Universal Driver</i> (Type 4 driver) JAR files, included with the DB2 for Linux, UNIX and Windows product installation, to the local computer where the Summarization and Pruning Agent is installed. You can copy the files to any directory that the user that the Summarization and Pruning Agent process runs as has access to.	<p>The Type 4 driver file names and locations are as follows:</p> <pre>db2installdir/java/db2jcc.jar db2installdir/java/db2jcc_license_cu.jar</pre> <p>where <i>db2installdir</i> is the directory where DB2 for Linux, UNIX and Windows was installed. The default DB2 for Linux, UNIX and Windows Version 9 installation directory is as follows:</p> <ul style="list-style-type: none"> • On Windows: C:\Program Files\IBM\SQLLIB • On AIX: /usr/opt/db2_09_01 • On Linux or Solaris: /opt/IBM/db2/V9.1
<p>Configure the Summarization and Pruning Agent.</p> <p>When you configure the Summarization and Pruning Agent, you configure the connection to the Tivoli Data Warehouse and you specify settings that control the operation of the Summarization and Pruning Agent.</p> <p>Perform this procedure whether or not the Summarization and Pruning Agent and the warehouse database are installed on the same computer.</p>	<p>“Configuring the Summarization and Pruning Agent (JDBC connection)” on page 643</p>
Configure the Summarization and Pruning Agent to connect to the Tivoli Enterprise Portal Server. Perform this procedure whether or not the Summarization and Pruning Agent and the warehouse database are installed on the same computer.	<p>See Step 9 of “Configuring the Summarization and Pruning Agent (JDBC connection)” on page 643.</p>
<p>Configure history collection.</p> <p>When you configure history collection, you specify settings for how often to collect, aggregate, and prune data for individual monitoring agents and attribute groups. Configure history collection from the Tivoli Enterprise Portal.</p>	<p>See the <i>IBM Tivoli Monitoring Administrator's Guide</i> for instructions on how to configure history collection.</p>
Start the Summarization and Pruning Agent.	<p>“Starting the Summarization and Pruning Agent” on page 656</p>

Step 5: Install and configure communications for the Tivoli Performance Analyzer

You can install Tivoli Performance Analyzer to a server which also has other Tivoli Monitoring components installed, or you can install it to a separate machine. The installation procedure is similar to that for monitoring agents.

Complete the tasks described in the following table, in the order listed, to install and configure Tivoli Performance Analyzer.

Table 91. Tasks for installing and configuring communications for the Tivoli Performance Analyzer

Task	Procedure
<p>The installation procedure for Windows includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring server. On Linux or AIX, this step is performed in a separate configuration procedure (Configuring the monitoring agent) and an X11 GUI is required to configure the agent. Alternatively, you can run the following command to utilize an X terminal emulation program (such as Cygwin) that is running on another computer:</p> <pre>export DISPLAY=my_windows_pc_IP_addr:0.0</pre> <p>where my_windows_pc_IP_addr is the IP address of a computer that is running an X terminal emulation program. See the information opposite. Be sure to perform all referenced installation and configuration procedures.</p> <p>Note: for sites setting up autonomous operation, the installation procedure includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring Server.</p>	<p>To install Tivoli Performance Analyzer on Windows, complete the procedure “Windows: Installing a monitoring agent” on page 262. To install Tivoli Performance Analyzer agent on Linux or AIX, complete the procedure “Linux or UNIX: Installing a monitoring agent” on page 267, including the following subsections:</p> <ul style="list-style-type: none"> • Installing the monitoring agent • Configuring the monitoring agent • Changing the file permissions for agents <p>Do not complete the procedure for starting the agent.</p>
On the computer where Tivoli Performance Analyzer is installed, configure an ODBC data source for the data warehouse.	“Configuring an ODBC data source for a DB2 data warehouse” on page 551
If the data warehouse is located on a remote computer, copy the JDBC Universal Driver (Type 4 driver) JAR files, included with the DB2 for Linux, UNIX, and Windows product installation, to the local computer where Tivoli Performance Analyzer is installed. You can copy the files to any directory on the local computer.	<p>The Type 4 driver file names and locations are the following:</p> <pre>db2installdir/java/db2jcc.jar db2installdir/java/db2jcc_license_cu.jar</pre> <p>where db2installdir is the directory where DB2 for Linux, UNIX, and Windows was installed. The default DB2 for Linux, UNIX, and Windows Version 9 installation directory is as follows:</p> <ul style="list-style-type: none"> • for AIX: /usr/opt/db2_09_01 • for Linux: /opt/IBM/db2/V9.1
Configure Tivoli Performance Analyzer to connect to the data warehouse.	For Tivoli Performance Analyzer on Windows, see “Configuring Tivoli Performance Analyzer on Windows (ODBC connection)” on page 565. For a Tivoli Performance Analyzer on Linux or AIX, see “Configuring Tivoli Performance Analyzer on Linux or UNIX (JDBC connection)” on page 566.
Start Tivoli Performance Analyzer.	“Starting Tivoli Performance Analyzer” on page 568

Configuring an ODBC data source for a DB2 data warehouse

A DB2 for Linux, UNIX and Windows client on Windows requires an ODBC connection to the data warehouse.

Before you begin

This procedure uses default values for the data source name, warehouse alias, and warehouse user ID. (Default values are used in configuration procedures for warehousing components.) Substitute different values if you do not want to use the default values.

About this task

Complete the following procedure to set up an ODBC connection for Tivoli Performance Analyzer on Windows to a local or remote Tivoli Data Warehouse.

Procedure

1. On the computer where the Tivoli Performance Analyzer is installed, open the Control Panel.
2. Click **Administrative Tools** → **Data Sources (ODBC)**
3. Click **Add** in the **System DSN** tab in the ODBC Data Source Administrator window.
4. Select **IBM DB2 ODBC DRIVER** from the list.
5. Click **Finish**.
6. In the ODBC DB2 Driver - Add window, perform the following steps:
 - a. Enter ITM Warehouse in **Data source name**.
 - b. Enter Warehous in **Database Alias**.

If the Tivoli Data Warehouse is located on a remote computer, ensure that the database alias matches the alias that you used when cataloging the remote data warehouse. See “Cataloging a remote data warehouse” on page 550.

If local, ensure that the database alias matches the name used for the warehouse database.
 - c. Click **OK**.
7. Test the ODBC database connection before continuing:
 - a. In the ODBC Data Source Administrator window, select **ITM Warehouse**.
 - b. Click **Configure**.
 - c. In the CLI/ODBC Settings - ITM Warehouse window, you see the data source name, **ITM Warehouse**.
 - d. Enter ITMUser for the **User ID**.
 - e. Type a password for the user in the **Password** field. The default password is `itmpswd1`.
 - f. Click **Connect**.
 - g. A **Connection test successful** message is displayed.
 - h. Click **OK**.
 - i. Click **OK** to close the window.

Configuring Tivoli Performance Analyzer on Windows (ODBC connection)

Use this procedure to configure Tivoli Performance Analyzer on Windows to connect to a DB2 for Linux, UNIX, and Windows data warehouse.

Procedure

1. Log on to the Windows system where Tivoli Performance Analyzer is installed and begin the configuration:
 - a. Click Start → Programs → IBM Tivoli Monitoring → Manage Tivoli Monitoring Services. The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click Performance Analyzer and click **Reconfigure**.
 - c. Click **OK** on the message regarding connection to a hub monitoring server.
2. The next two windows (entitled Performance Analyzer: Agent Advanced Configuration) contain the settings for the connection between Tivoli Performance Analyzer and the hub monitoring server. These settings were specified when Tivoli Performance Analyzer was installed. Click **OK** on each window to accept the settings.
3. Click **Yes** on the message asking if you want to configure the ODBC data source.

4. Select **ODBC** from the list of selectable agent database connection types.
5. Set your Database Type to DB2.
6. Specify the Data Source Name - Agent ODBC DSN (ITM Warehouse is the default).

Note: Tivoli Performance Analyzer does not create this DSN - it must already exist. If you are installing the agent on the same machine where TEP Server is installed, you can use the existing data source created by Tivoli Monitoring. Otherwise, you must create a new System DSN manually, prior to re-configuring Tivoli Performance Analyzer.

On 64-bit versions of Windows, data sources created by the default ODBC Data Source Administrator applet available from the Control Panel are not available for 32-bit applications. Therefore you must use the 32-bit version of the ODBC Data Source Administrator applet from <WINDOWS>\SysWOW64\odbcad32.exe.

Values for the data source name, database name, and database user ID and password must match the values that you used when configuring an ODBC connection for Tivoli Performance Analyzer. For more information, see “Configuring an ODBC data source for a DB2 data warehouse” on page 564.

7. Type the Username and Password. The entries in these fields are used to connect to the Tivoli Data Warehouse and are the same credentials as those used by the Tivoli Enterprise Portal Server, the Warehouse Proxy Agent, the Summarization and Pruning Agent, and the Performance Analyzer to communicate with the Tivoli Data Warehouse.
8. Click **Next** to proceed to the Advanced Configuration window.
9. You can enable Advanced Configuration to specify Tivoli Data Warehouse schema and The Tivoli Data Warehouse database schema. If you do not select **Enable advanced configuration** these options are greyed out.
10. You can also choose whether you want the agent to Initialize PA tables and OS domain tasks.

Note: Setting Initialize PA tables to YES will remove and recreate all previously created tables deleting all user tasks and reverting each OS task to its default.

11. Use the **Bypass connection tests** option to finish the configuration without running connection tests.
12. Click **OK** to finish the configuration process.

Table 92. Configuration information for the Tivoli Data Warehouse database on DB2 for Linux, UNIX and Windows

Field	Default value	Description
ODBC DSN	ITM Warehouse	The name of the data source.
Username	ITMUser	The name of the Windows OS user that the Tivoli Performance Analyzer will use to access the Tivoli Data Warehouse database.
Password	itmpswd1	The password for the Windows OS user. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Test Connection		Test the connection to the Tivoli Data Warehouse Database based on the completed fields above: ODBC DSN, Username, and Password.

Configuring Tivoli Performance Analyzer on Linux or UNIX (JDBC connection)

Use this procedure to configure Tivoli Performance Analyzer on Linux or UNIX to connect to a DB2 for Linux, UNIX, and Windows data warehouse on any operating system.

Procedure

1. To begin the configuration, log on to the computer where Tivoli Performance Analyzer is installed.
 - a. Change to the `install_dir/bin` directory and run the following command:

```
./itmcmd manage [-h install_dir]
```

where `install_dir` is the installation directory for IBM Tivoli Monitoring. The default installation directory is `/opt/IBM/ITM`. The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Performance Analyzer** and click **Configure**. The Configure Tivoli Performance Analyzer window is displayed.
2. Set the Database Type to DB2.
3. Type the username and the password. The entries in these fields are used to connect to the Tivoli Data Warehouse.
4. Review all the defaults in the Agent Configuration window and change as required.
 - a. If the Tivoli Data Warehouse is installed on a remote computer, specify the host name of the remote computer instead of `localhost`.
 - b. Change the port number if necessary (the default port number for DB2 is 50000).
 - c. If the name of the Tivoli Data Warehouse database is not `WAREHOUS`, replace `WAREHOUS` with the actual name. (See “Creating the warehouse database on DB2 for Linux, UNIX and Windows” on page 544.)

Note: when specifying the DB2 for Linux, UNIX, and Windows database name on Linux and UNIX, case is ignored. In other words, it makes no difference whether you provide the database name in lowercase or uppercase.

5. Specify the JDBC Driver. The default driver name for DB2 is `com.ibm.db2.jcc.DB2Driver`.
6. Specify the JDBC Driver Path, which should be provided as a list of JAR files with the full path separated by ":".

Note: The driver files names for DB2 for Linux, UNIX, and Windows are:

```
db2jcc.jar  
db2jcc_license_cu.jar
```

Fast path: You can use the Browse button to specify the path. In such a case a file list is added at the end of the JDBC Driver Path text field, separated from the existing content by a path separator.

Attention: For compatibility with version 6.1.1 it is also possible to provide a single directory name containing all the JAR files which should be used, although this solution is not recommended. Ensure that the directory does not contain the `db2jcc4.jar` file. Performance Analyzer only supports Java version 1.5 and requires the `db2jcc.jar` file. Including `db2jcc4.jar` in the classpath may prevent the agent from being successfully configured.

7. You can use the **Test connection** button to check whether the connection can be initiated.
8. Click **Next** to proceed to the Advanced Configuration window.
 - a. You can enable Advanced Configuration to specify TDW Schema and Configuration schema. If you do not select Enable advanced configuration, all of these options are greyed out.
 - b. You can also choose whether you want the agent to initialize PA tables.

Note: Setting Initialize PA tables to YES will remove and recreate all previously created tables deleting all user tasks and reverting each OS task to its default.

- c. Use the **Bypass connection tests** option to finish the configuration without running connection tests.
9. Click **Save** to save your settings and close the window.

Starting Tivoli Performance Analyzer

This section contains instructions for starting the Tivoli Performance Analyzer.

To start Tivoli Performance Analyzer from the Manage Tivoli Enterprise Monitoring Services window, right-click Tivoli Performance Analyzer and select Start. To start the Tivoli Performance Analyzer agent from the command-line, run the following command from the bin directory of the IBM Tivoli Monitoring installation directory. The default installation directory is /opt/IBM/ITM.

```
./itmcmd agent start pa
```

where pa is the product code for Tivoli Performance Analyzer agent.

Chapter 24. Tivoli Data Warehouse solution using DB2 on z/OS

Use the information and instructions in this chapter to implement a Tivoli Data Warehouse solution using mainframe-based DB2 running on z/OS as your warehouse database.

The following table lists the goals for creating a DB2 on z/OS solution.

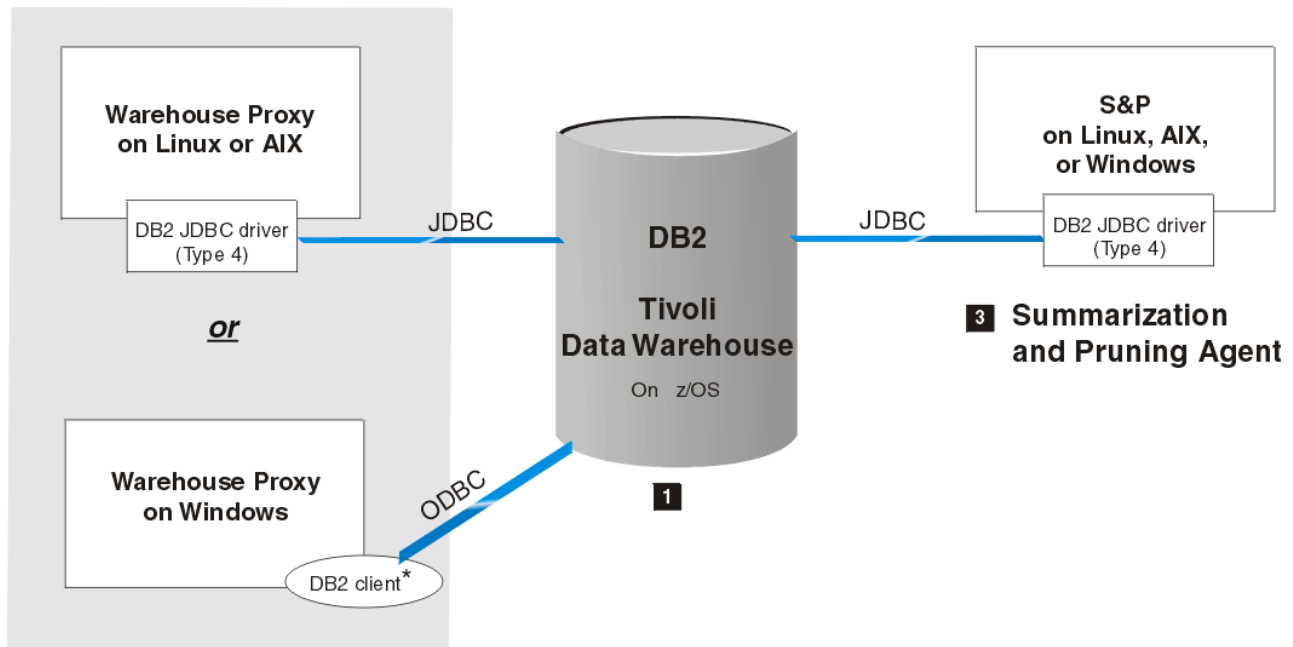
Table 93. Goals for creating a Tivoli Data Warehouse solution using DB2 on z/OS

Goal	Where to find information
Review your options, specific to a DB2 on z/OS solution, for operating system platforms and communications between warehousing components.	"Supported components"
Install prerequisite software before implementing your Tivoli Data Warehouse solution.	"Prerequisite installation" on page 571
Understand how to use the instructions for implementing your Tivoli Data Warehouse solution.	"Implementing a Tivoli Data Warehouse solution using DB2 on z/OS" on page 571
Complete the steps for implementing your Tivoli Data Warehouse solution using DB2 on z/OS for the data warehouse.	"Step 1: Connect the Warehouse Proxy node to your DB2 on z/OS database" on page 573 "Step 2: Configure the Tivoli Data Warehouse agents" on page 585

Supported components

The diagram in this section summarizes the supported operating system platforms for the various warehousing components, the supported database products, and the connections between components.

Figure 106 on page 570 presents the IBM Tivoli Monitoring environment when implementing a Tivoli Data Warehouse solution using DB2 on z/OS as the warehouse database. For more specific information about supported operating systems and database products, including product names and versions, see Chapter 3, "Hardware and software requirements," on page 75.



2 Warehouse Proxy Agent

Figure 106. Tivoli Data Warehouse solution using DB2 on z/OS

Note: An asterisk (*) next to a database client indicates that you must manually install the client if it does not already exist.

In the following discussion, numbered product components correspond to the numbers on the diagram.

1 Tivoli Data Warehouse on DB2 on z/OS

A Tivoli Data Warehouse repository on DB2 on z/OS can be accessed by any ITM-supported platform that can run the Warehouse Proxy Agent—Windows, Linux, or AIX—as well as the DB2 Connect software.

2 Warehouse Proxy Agent

A Warehouse Proxy Agent on Linux or AIX communicates with the warehouse database through a JDBC connection. Install a Type 4 driver (DB2 on z/OS JDBC Universal Driver) on the computer where the Warehouse Proxy Agent is located.

A Warehouse Proxy Agent on Windows communicates with the warehouse database through an ODBC connection. The ODBC driver is included with the DB2 on z/OS client. You must install a DB2 client on the Windows computer where the Warehouse Proxy Agent is located, and then catalog the remote node and database on the local computer.

Note: If you install a version 9 DB2 client, you also must install DB2 Connect Server Edition to connect to a DB2 on z/OS data server.

3 Summarization and Pruning Agent

The Summarization and Pruning Agent communicates with the warehouse database through a JDBC connection from any supported operating system. Install a DB2 on z/OS Type 4 JDBC driver (*DB2 on z/OS JDBC Universal Driver*) on the computer where the Summarization and Pruning Agent is located.

Prerequisite installation

Before you implement your Tivoli Data Warehouse solution, complete one or more hub installations, including the warehousing components (see the appropriate chapter within this section for the necessary installation instructions).

Include the following components in each hub installation:

- The hub Tivoli Enterprise Monitoring Server
- (Optional) One or more remote monitoring servers
- The Tivoli Enterprise Portal Server, on either Windows, Linux, or UNIX
- An IBM DB2 on z/OS server on the computer where you will create the Tivoli Data Warehouse database. (The Tivoli Data Warehouse database can be shared in a multi-hub installation or dedicated to a single hub.)
- DB2 Connect Server Edition
- (Optional) A portal desktop client
- (Optional) Monitoring agents and the application support for the monitoring agents
- The Warehouse Proxy Agent and the Summarization and Pruning Agent
- (Optional) The Tivoli Performance Analyzer
- (Optional) Language packs for all languages other than English

See Table 94 for related information:

Table 94. Information topics related to installation of prerequisite software for a Tivoli Data Warehouse solution

Topic	Where to find information
Single and multiple hub installations	To understand the terminology related to single and multiple hub installations, see “Locating and sizing the hub Tivoli Enterprise Monitoring Server” on page 35.
Installation procedures for prerequisite components, excluding the Warehouse Proxy Agent and the Summarization and Pruning Agent	The detailed instructions for installing the prerequisite components are described in Chapter 10, “Installing IBM Tivoli Monitoring,” on page 193. See your database documentations for instructions on how to install a supported database server.
Supported RDBMS versions	For specific information about the supported database platforms for the portal server database and the Tivoli Data Warehouse, see Chapter 3, “Hardware and software requirements,” on page 75.

Implementing a Tivoli Data Warehouse solution using DB2 on z/OS

Use the instructions in the remainder of this chapter to implement a Tivoli Data Warehouse solution using DB2 on z/OS version 9.1 (or subsequent) as your data warehouse.

Requirements:

You must set permissions for both the Warehouse Proxy Agent and the Summarization and Pruning Agent when connecting to a DB2 database on z/OS. In the commands that follow, USER1 is in place of the user you are connecting to the Warehouse Proxy Agent or the Summarization and Pruning Agent.

To configure the Warehouse Proxy Agent run the following command:

```
db2 "GRANT CREATEDBA TO USER1"
```

To grant the permission that is required to start the Warehouse Proxy Agent run the following command:

```
db2 "GRANT SELECT ON SYSIBM.SYSTABLES TO USER1"
```

Both of the following access permissions must be granted for the Summarization and Pruning Agent:

```
db2 "GRANT SELECT ON SYSIBM.SYSROUTINES TO USER1"
db2 "GRANT SELECT ON SYSIBM.SYSTABLESPACE TO USER1"
```

The Warehouse Proxy Agent and the Summarization and Pruning Agent both use the implicit table creation option when connected to a DB2 database on z/OS. Both agents create tables without specifying a table space or a database in the IN clause of a CREATE statement.

- DB2 on z/OS version 9 creates an implicit database each time a table is created using a name in the range DSN00001 to DSN60000. The characteristics of implicitly created databases are shown in Table 95.

Table 95. Characteristics of implicitly created DB2 on z/OS database

Field	Value
Name	DSNxxxxx, where xxxxx is a number from 00001 to 60000
BUFFERPOOL	BP0, BP8K0, BP16K0, BP32K ¹ Default values. Changeable through DSNZPARM update
INDEXBP	IDXBPOOL setting in DSNZPARM
STOGROUP	SYSDEFLT
value in column IMPLICIT of SYSIBM.SYSDATABASE	'Y'
ENCODING_SCHEME	DEFAULT is the DSNHDECP setting (see below)
SBCS_CCSID	DEFAULT is the DSNHDECP setting (see below)
DBCS_CCSID	DEFAULT is the DSNHDECP setting (see below)
MIXEC_CCSID	DEFAULT is the DSNHDECP setting (see below)

Note:

1. DB2 on z/OS chooses a specific buffer pool during creation of the implicit object, depending on the record size. When the maximum record size reaches approximately 90% of the capacity of the smaller page size, DB2 chooses the next larger page size, as shown in Table 96.

Table 96. Maximum DB2 on z/OS page sizes

Name	Page Size
BP0	4K
BP8KO	8K
BP16KO	16KB
BP32K	32KB

2. To ensure databases can be created implicitly, DB2 on z/OS's CREATE IMPLICIT DATABASES installation parameter must be set to YES.
- DB2 on z/OS v9 creates implicit table spaces when implicitly creating a table.
 - The Summarization and Pruning Agent creates functions that Tivoli Enterprise Portal user can take advantage of when creating custom history queries. To let the agent create those functions, a default WLM (Workload Manager) should be created.
 - Before beginning this procedure, gather the following information about your target DB2 on z/OS database:

Table 97. Required parameters for accessing the DB2 on z/OS database

DB2 on z/OS parameter	Your value
Database name	
Port number	

Table 97. Required parameters for accessing the DB2 on z/OS database (continued)

DB2 on z/OS parameter	Your value
DB2 userid	
DB2 password	
Fully qualified host name	

Solution steps

To implement your Tivoli Data Warehouse solution using DB2 on z/OS, complete the major steps described in the remaining sections of this chapter, in the order listed.

1. Connect the Warehouse Proxy node to your DB2 on z/OS database.
2. Configure the Tivoli Data Warehouse agents.

To implement your solution successfully:

- Perform the tasks in the order listed.
- Do not skip a task and move forward to the procedures that follow it.

Step 1: Connect the Warehouse Proxy node to your DB2 on z/OS database

On the IBM Tivoli Monitoring node running the Warehouse Proxy Agent, invoke the DB2 Client Configuration Assistant from the DB2 Control Center.

The Client Configuration Assistant, shown in Figure 107 on page 574, opens.

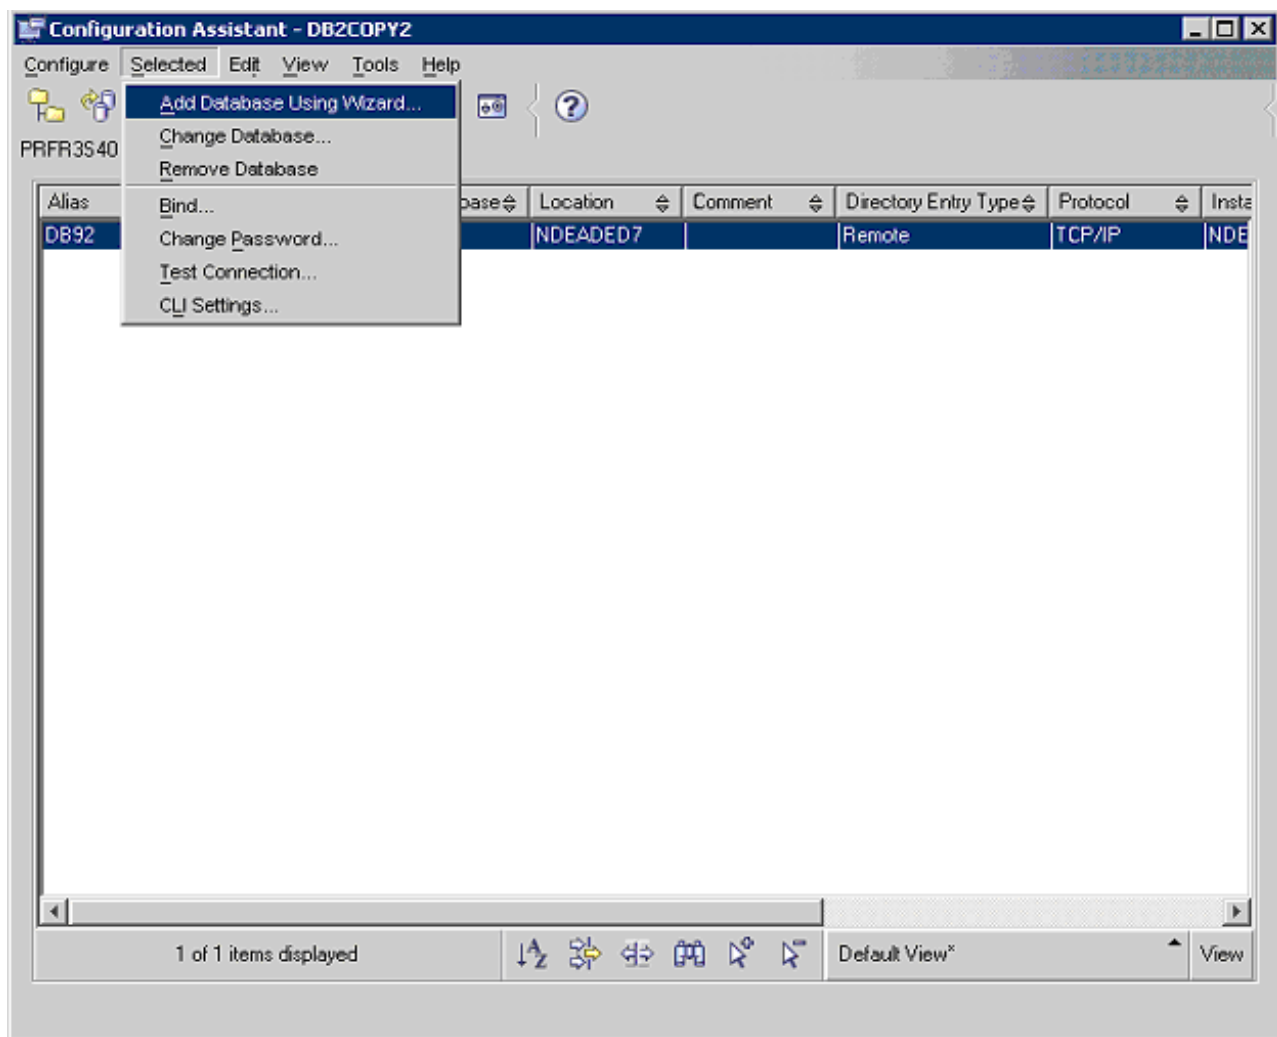


Figure 107. DB2 Client Configuration Assistant screen

Start defining the database connection

From the **Selected** pulldown menu, select the **Add Database Using Wizard** menu option to set up the connection to your DB2 on z/OS database.

The Add Database Wizard notebook opens with the **Source** tab active, as shown in Figure 108 on page 575.

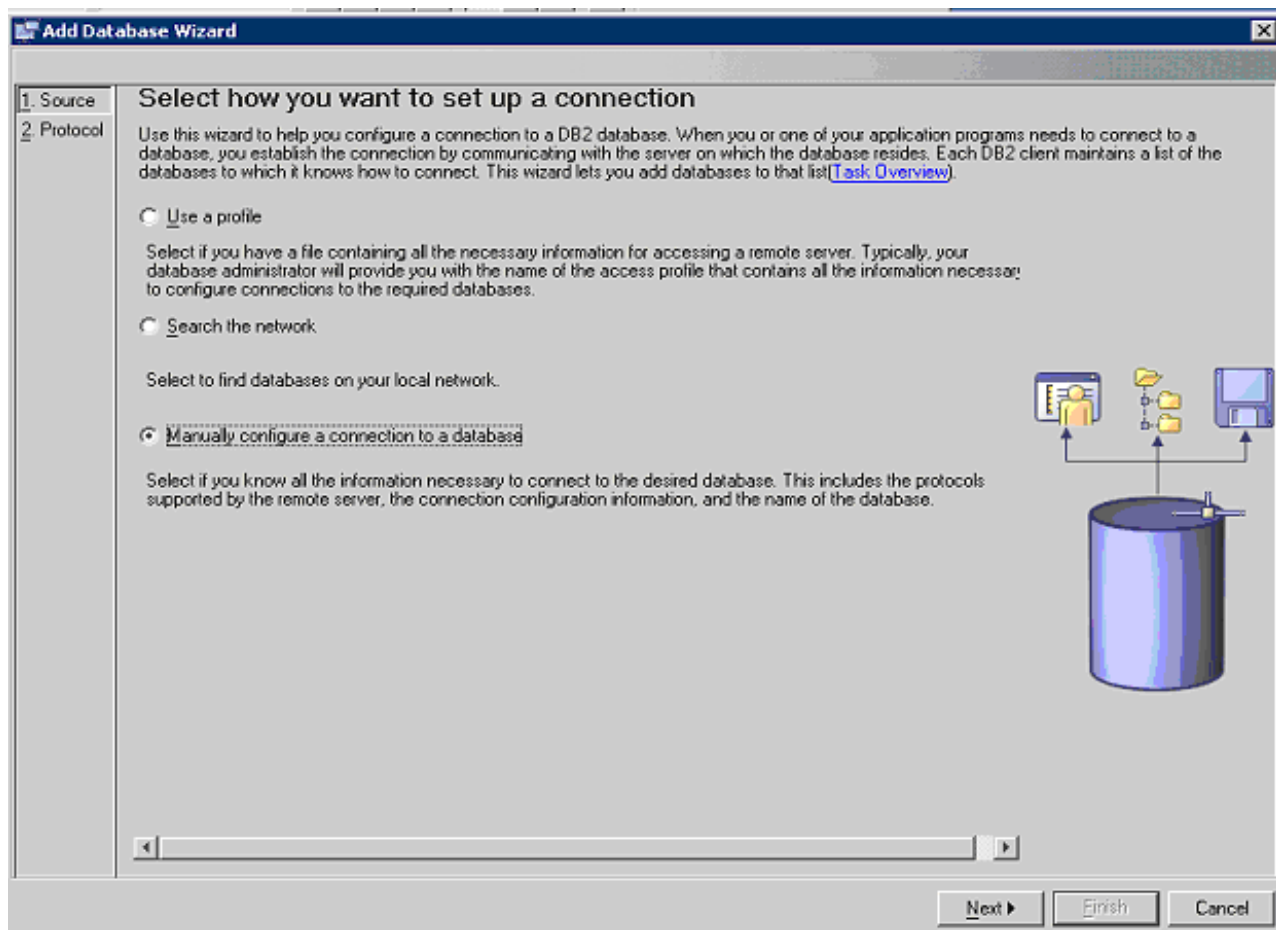


Figure 108. DB2 Add Database Wizard notebook, Source tab

Activate the **Manually configure a connection to a database** radio button; then click **Next**.

Define the communications protocol

Complete the fields in this section to define your communications protocol.

The Add Database Wizard notebook reappears with the **Protocol** tab active, as shown in Figure 109 on page 576.




Figure 109. DB2 Add Database Wizard notebook, Protocol tab

On the Protocol notebook page:

1. Activate the **TCP/IP** radio button.
2. Select the **The database physically resides on a host or AS/400 system** option.
3. Activate the **Connect directly to the server** radio button.
4. Click **Next**.

Define the TCP/IP communications parameters

Complete the fields in this section to specify your TCP/IP communication parameters.

The Add Database Wizard notebook reappears with the **TCP/IP** tab active, as shown in Figure 110 on page 577.

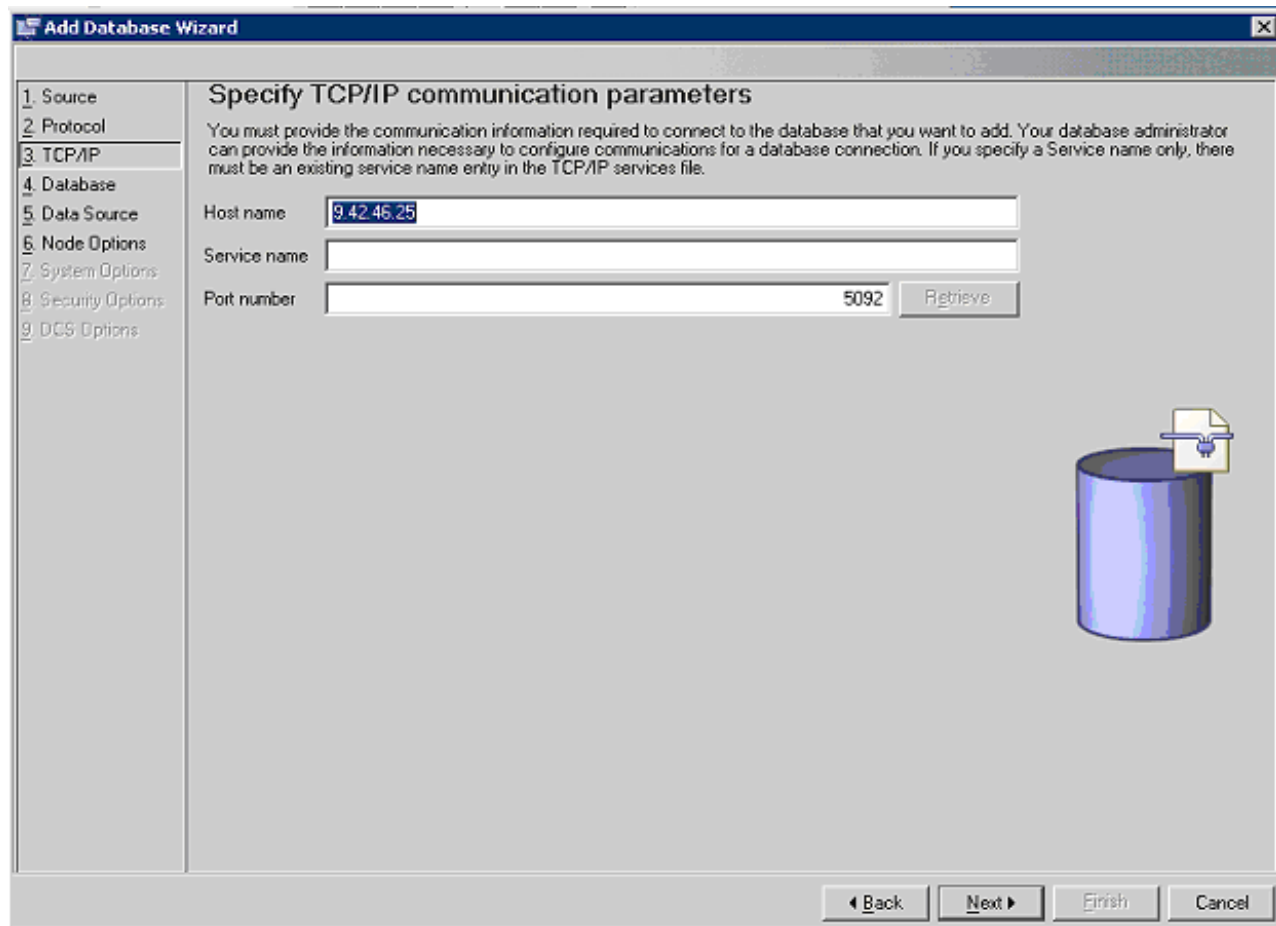


Figure 110. DB2 Add Database Wizard notebook, TCP/IP tab

On the TCP/IP notebook page:

1. Enter the fully qualified IP **Host name**.
2. Specify the **Port number**. The default port for z/OS is 446.
3. Click **Next**.

Identify the DB2 on z/OS database

Specify the name of the database to which you want to connect.

The **Database** tab opens, as shown in Figure 111 on page 578.




Figure 111. DB2 Add Database Wizard notebook, Database tab

On the Database notebook page:

1. Enter the fully qualified **Database name**. For DB2 for Linux, UNIX, and Windows, this is the database name; however, with DB2 on z/OS, this is the subsystem name.
2. Enter the DB2 on z/OS **Database alias**. This field is required.

Note: The alias is limited to 8 characters. Do not create duplicate aliases across your Tivoli Monitoring systems (such as WAREHOUSE).

3. Click **Next**.

Register the database as an ODBC data source

This section illustrates the Data Source tab of the DB2 Add Database Wizard.

The **Data Source** tab opens, as shown in Figure 112 on page 579.

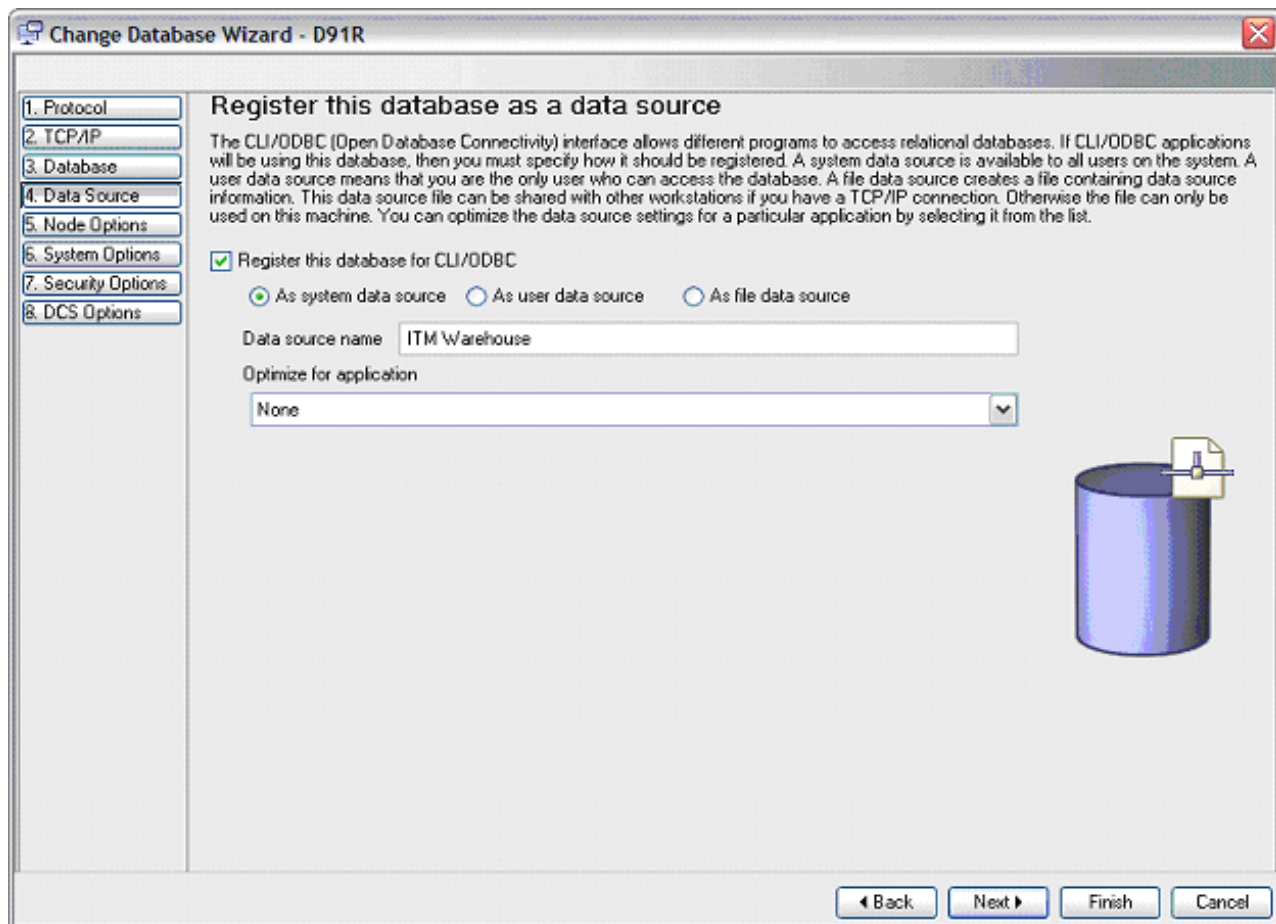


Figure 112. DB2 Add Database Wizard notebook, Data Source tab

On the Data Source notebook page, click **Next**.

1. Ensure the **Register this database for CLI/ODBC** option is selected.
2. Ensure the **As system data source** radio button is active.
3. Enter the fully qualified DB2 on z/OS **Data source name**. This is the name you will specify when configuring the Warehouse Proxy Agent in “Step 2: Configure the Tivoli Data Warehouse agents” on page 585, for example, ITM Warehouse.
4. Click **Next**.

Identify the z/OS server containing the DB2 on z/OS database

Specify the node options in the Add Database Wizard.

The **Node Options** tab opens, as shown in Figure 113 on page 580.

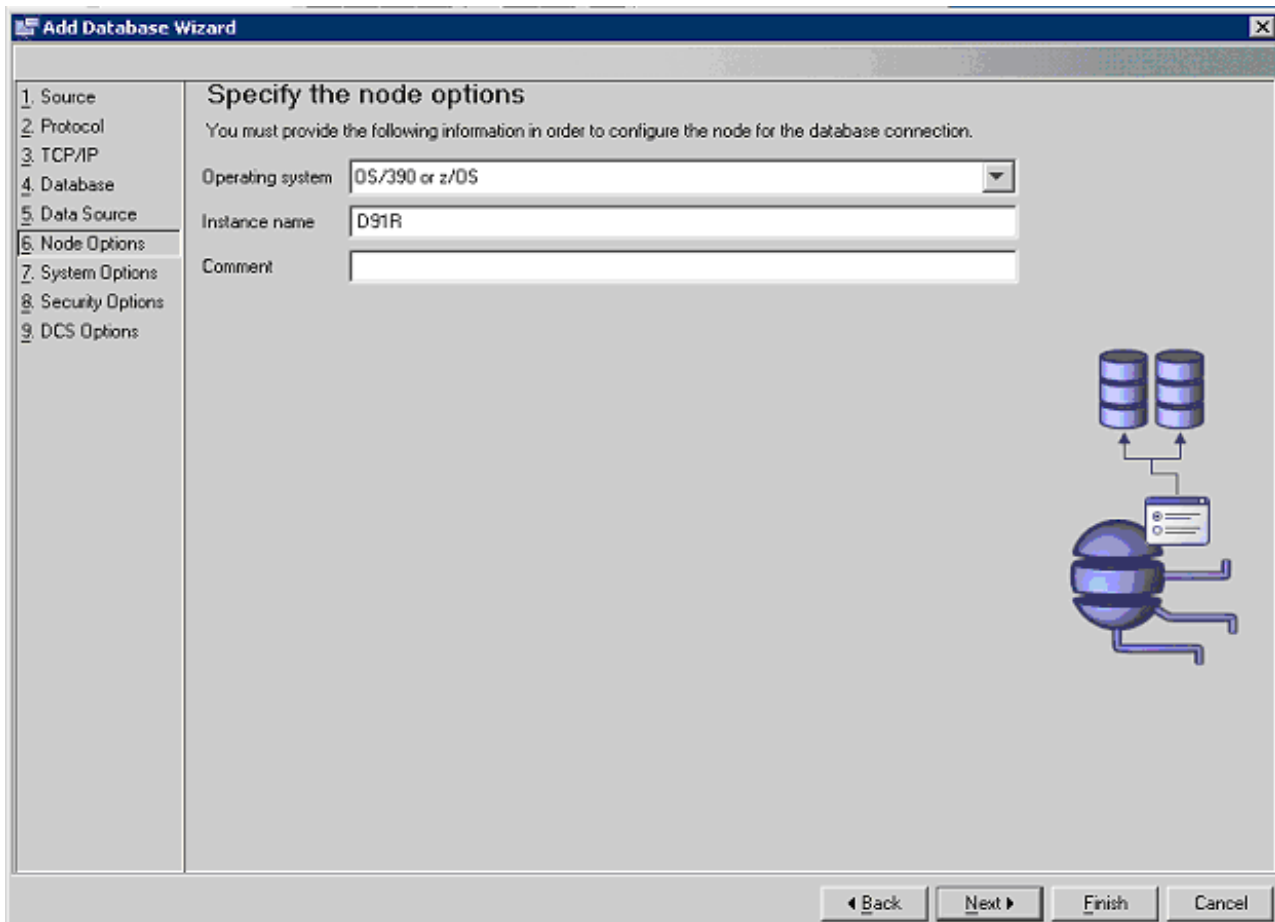


Figure 113. DB2 Add Database Wizard notebook, Node Options tab

On the Node Options notebook page:

1. Pull down the list of **Operating system** settings, and select **OS/390® or z/OS**.
2. For **Instance name**, specify your DB2 on z/OS instance name.
3. Click **Next**.

Define the DB2 on z/OS system options

On the System Options tab, provide the System name, Host name, and Operating system.

The **System Options** tab opens, as shown in Figure 114 on page 581.

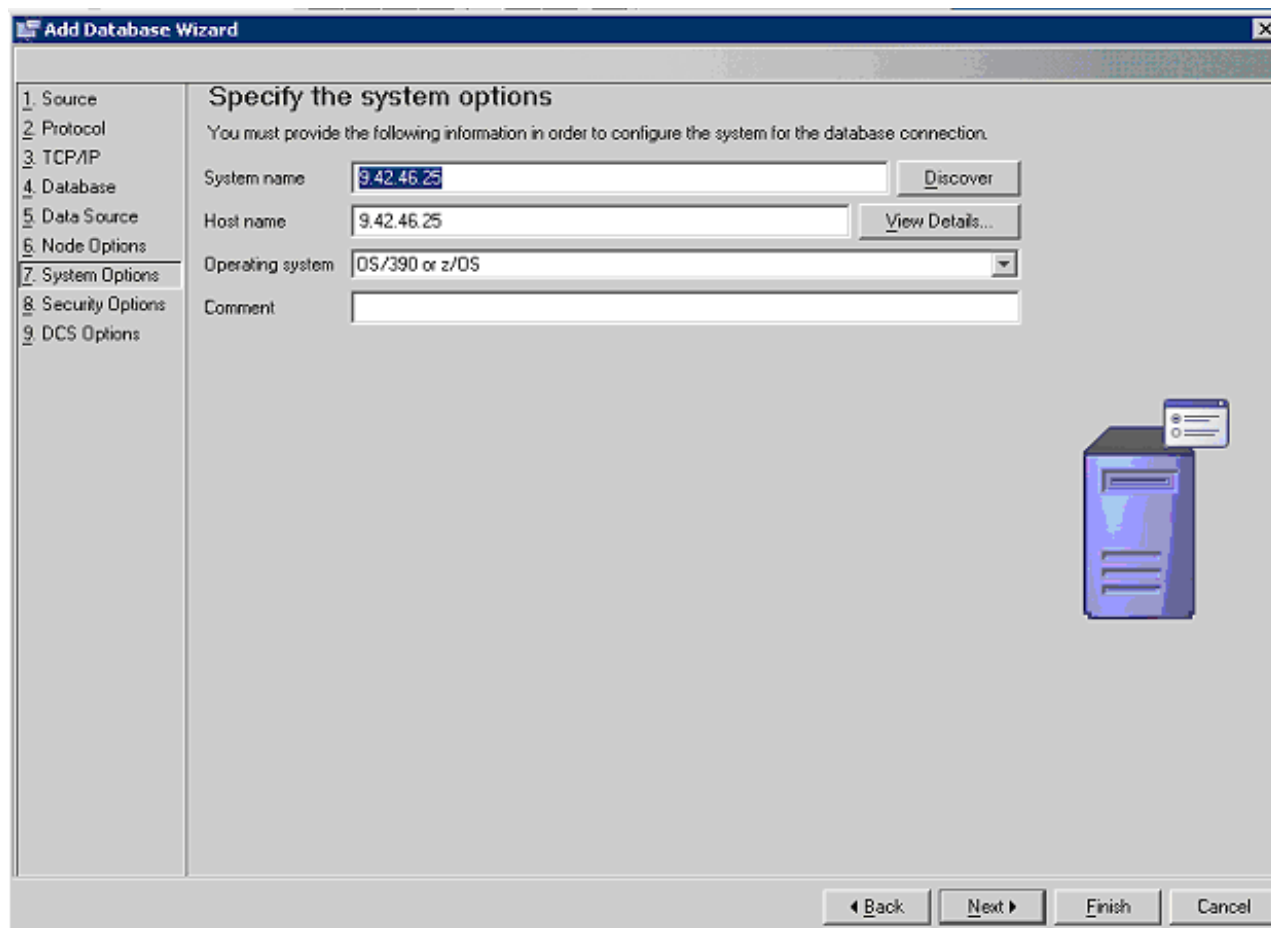


Figure 114. DB2 Add Database Wizard notebook, System Options tab

On the System Options notebook page:

1. For **System name**, specify the host name or IP address of the z/OS system where your DB2 on z/OS database is stored.
The **Host name** field gets filled in automatically with the same value.
2. Pull down the list of **Operating system** settings, and select **OS/390 or z/OS**.
3. Click **Next**.

Define the DB2 on z/OS security options

This section illustrates the Security Options tab of the DB2 Add Database Wizard.

The **Security Options** tab opens, as shown in Figure 115 on page 582.




Figure 115. DB2 Add Database Wizard notebook, Security Options tab

On the Security Options notebook page, activate the **Server authentication (SERVER)** radio button, and click **Next**.

Complete the DB2 on z/OS host connection

This section illustrates the DCS Options tab of the DB2 Add Database Wizard.

The **DCS Options** tab opens, as shown in Figure 116 on page 583.

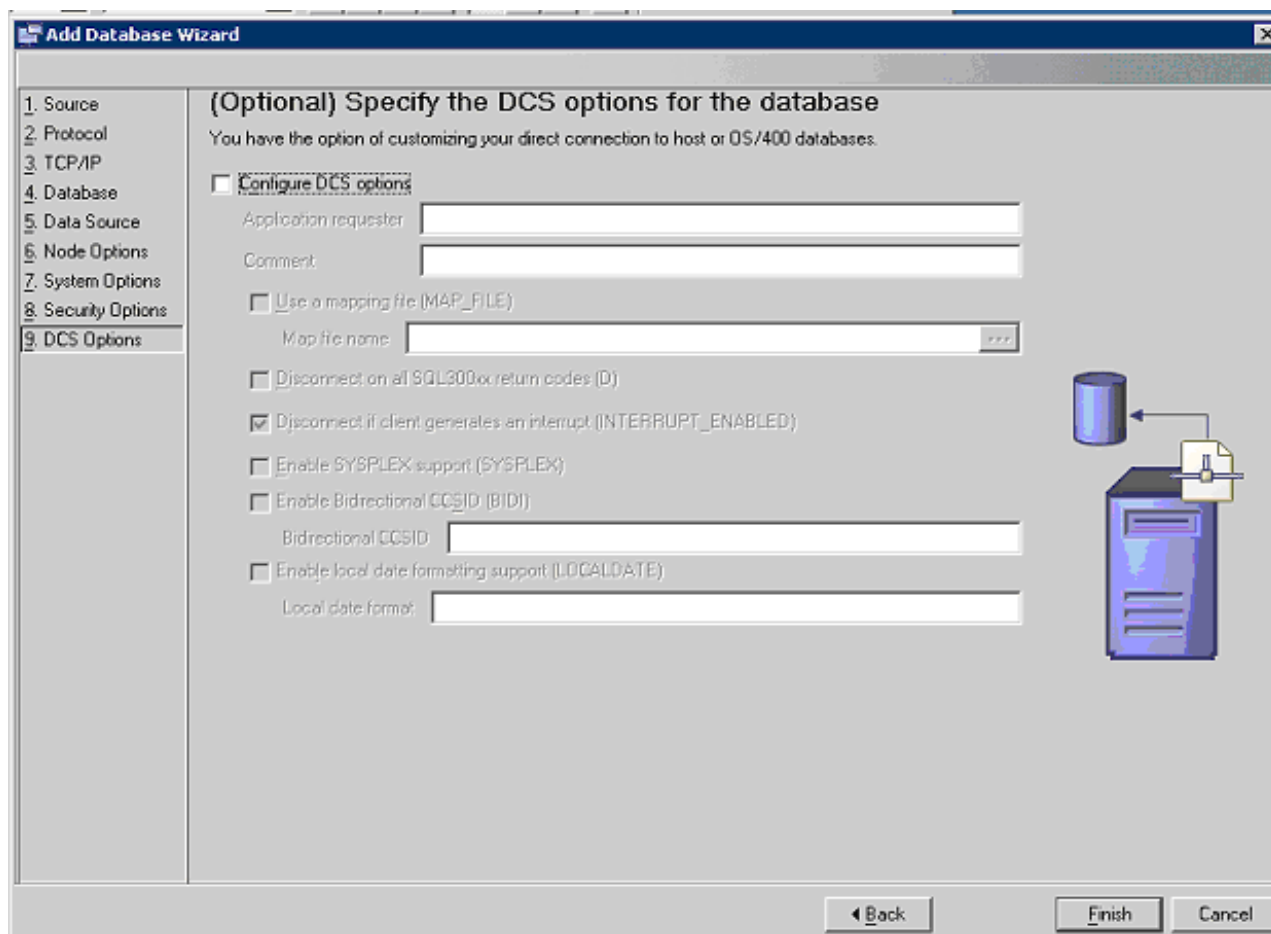


Figure 116. DB2 Add Database Wizard notebook, DCS Options tab

On the DCS Options notebook page, click **Finish**.

Verify that the connection can be made

If your database definition is successful, the Add Database Confirmation window is displayed.




Figure 117. Connection-confirmation screen

If this window is not displayed, press the **Back** button, verify each notebook page, and correct the information as necessary.

To test the connection to the remote DB2 on z/OS database, press the **Test Connection** button.

The **Connect to DB2 Database** screen, shown in Figure 118, opens.




Figure 118. Connect to DB2 Database screen

Enter the user ID and password for the remote DB2 on z/OS database, and press the **Test Connection** button.

If the database connection can be made, the confirmation screen shown in Figure 119 is displayed.




Figure 119. DB2 Connection Confirmation screen

Step 2: Configure the Tivoli Data Warehouse agents

To enable storage of IBM Tivoli Monitoring historical data in your Tivoli Data Warehouse database on z/OS, you need to perform the installation procedures for the Warehouse Proxy Agent and the Summarization and Pruning Agent.

Complete these two steps from Chapter 23, "Tivoli Data Warehouse solution using DB2 for Linux, UNIX, and Windows," on page 539:

- "Step 2: Install and configure communications for the Warehouse Proxy Agent" on page 548

Note: When specifying the DB2 on z/OS database name on Linux and UNIX, the correct case must be used.

- "Step 4: Install and configure communications for the Summarization and Pruning Agent" on page 562

Testing the database connection using the DB2 Control Center

This section illustrates using the DB2 Control Center to test the database connection.

Invoke the DB2 Control Center, shown in Figure 120 on page 586.

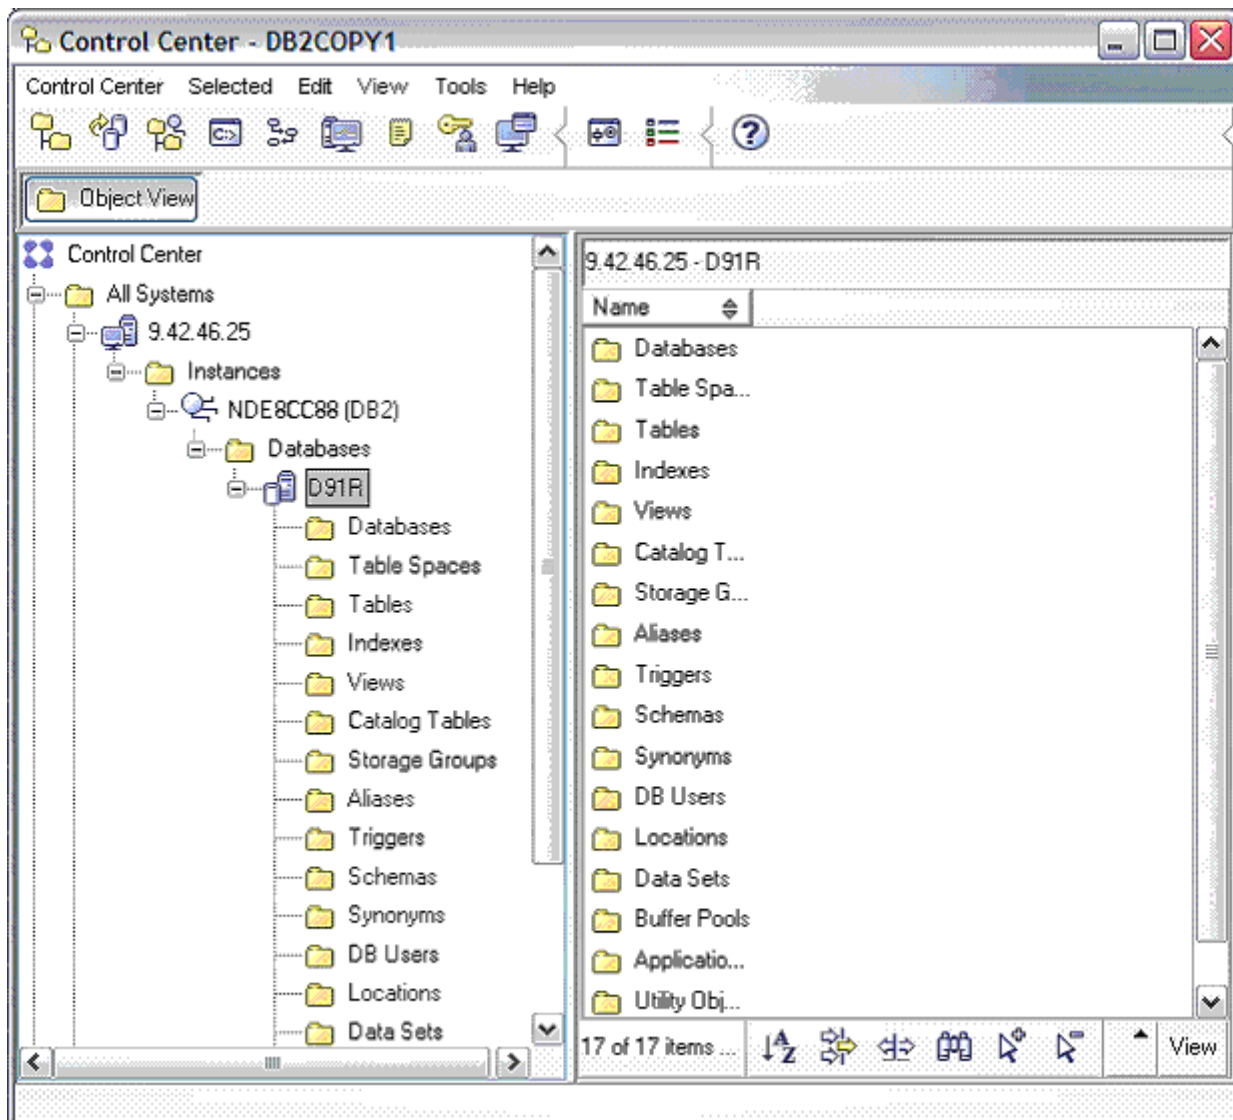


Figure 120. DB2 Control Center

Select the plus sign (+) to the left of the z/OS system that is running the DB2 on z/OS subsystem that owns your Tivoli Data Warehouse repository; then expand its list of databases. The list should include the database you're using.

To again check your database connection, right-click your database name, and select **Connect** from the pop-up menu, as shown in Figure 121 on page 587.




Figure 121. DB2 Control Center right-click action menu

The Connect to the *database* window, shown in Figure 122 on page 588, opens.



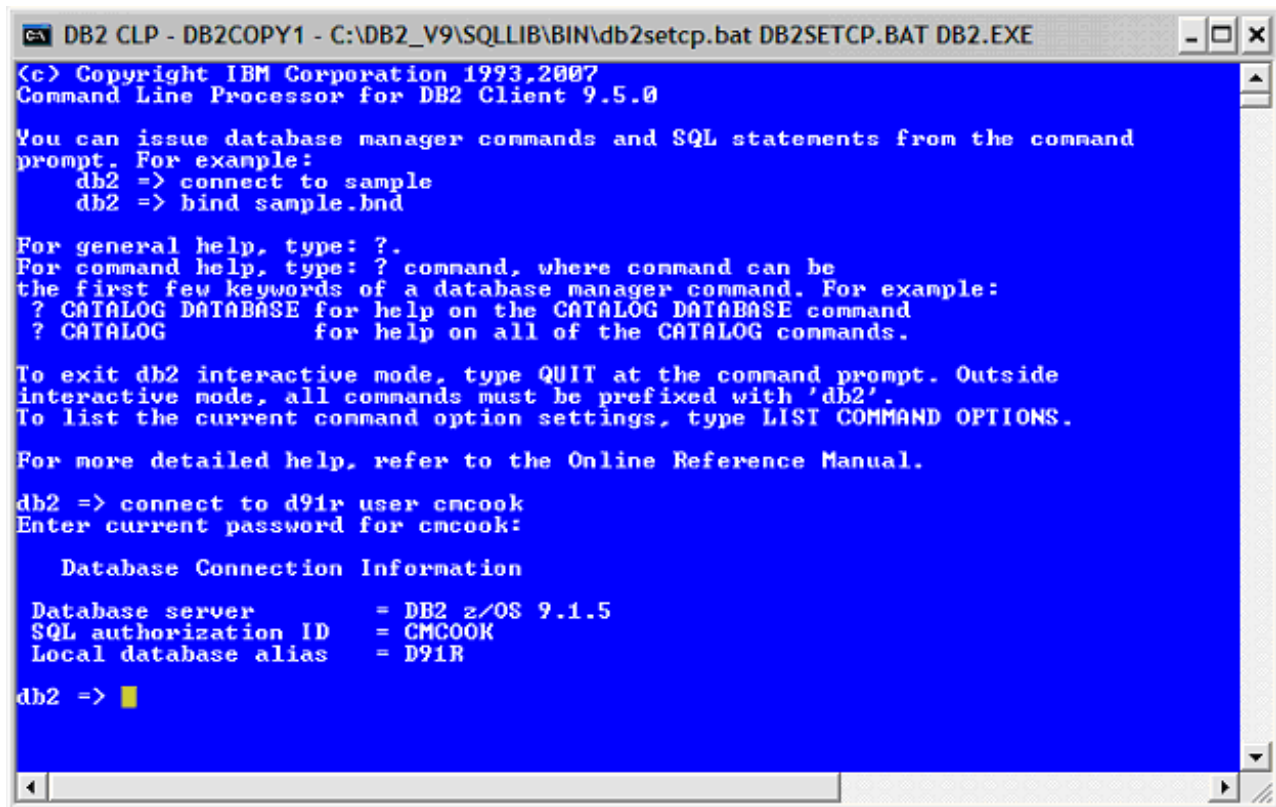
Figure 122. Connect to named database window

Enter the user ID and password for the remote DB2 on z/OS database, and press the **OK** button.

Testing the database connection using the DB2 command-line processor

This section illustrates using the DB2 command-line processor to test the database connection.

Invoke the DB2 command-line Processor, shown in Figure 123 on page 589.



```
<c> Copyright IBM Corporation 1993,2007
Command Line Processor for DB2 Client 9.5.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
    db2 => connect to sample
    db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
? CATALOG DATABASE for help on the CATALOG DATABASE command
? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => connect to d91r user cmcook
Enter current password for cmcook:

    Database Connection Information

Database server      = DB2 z/OS 9.1.5
SQL authorization ID = CMCOOK
Local database alias = D91R

db2 => █
```

Figure 123. DB2 command-line Processor window. Note that the password is hidden in this figure.

Enter the following command at the db2 prompt:

`connect to database_name user userid using password`

where *database_name* is the name you assigned to the DB2 on z/OS database that contains your data warehouse, and *userid* and *password* are the DB2 user ID and password required to access that database.

The database information is displayed, as shown in Figure 123.

Scenario: Connecting a Tivoli Data Warehouse solution using DB2 on z/OS to a Tivoli Enterprise Monitoring Server running z/OS

This scenario covers both Windows and UNIX/Linux machines. All agent configuration windows in this scenario should point to the hub Tivoli Enterprise Monitoring Server on z/OS.

On Windows:

A Warehouse Proxy Agent on Windows communicates with the warehouse database through an ODBC connection. The ODBC driver is included with the DB2 on z/OS client. You must install a DB2 client on the Windows computer where the Warehouse Proxy Agent is located, and then catalog the remote node and database on the local computer.

Note: If you install a version 9 DB2 client, you also must install DB2 Connect Server Edition on the workstation to connect to a DB2 on z/OS data server.

- You must first create an ODBC datasource on the Windows machine that will connect to the DB2 on z/OS database. You can do this using the DB2 Client Configuration Assistant tool, or through the

command-line. For instructions on using the DB2 Client Configuration Assistant tool, see “Step 1: Connect the Warehouse Proxy node to your DB2 on z/OS database” on page 573.

- Or run the following commands:

```
db2 catalog tcpip node DBNODE remote DBserverhostname server port number ostype OS390
db2 catalog dcs database db_name as db_name
db2 catalog db databasename-on server as alias-on client-databasename at node DBNODE authentication dcs
```

Use this procedure to reconfigure the Warehouse Proxy with the ODBC datasource you just created:

1. Log on to the Windows system where the Warehouse Proxy Agent is installed and begin the configuration:
 - a. Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Monitoring Services**.
The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Warehouse Proxy** and click **Configure Using Defaults**.
If the Warehouse Proxy is installed on the same computer as the portal server, click **Reconfigure**.
 - c. Click **OK** on the message regarding connection to a hub monitoring server.
2. The next two windows (entitled Warehouse Proxy: Agent Advanced Configuration) contain the settings for the connection between the Warehouse Proxy Agent and the hub monitoring server. These settings were specified when the Warehouse Proxy Agent was installed. Click **OK** on each window to accept the settings.
3. Click **Yes** on the message asking if you want to configure the ODBC data source.
4. Select **DB2** from the list of selectable databases and click **Next**

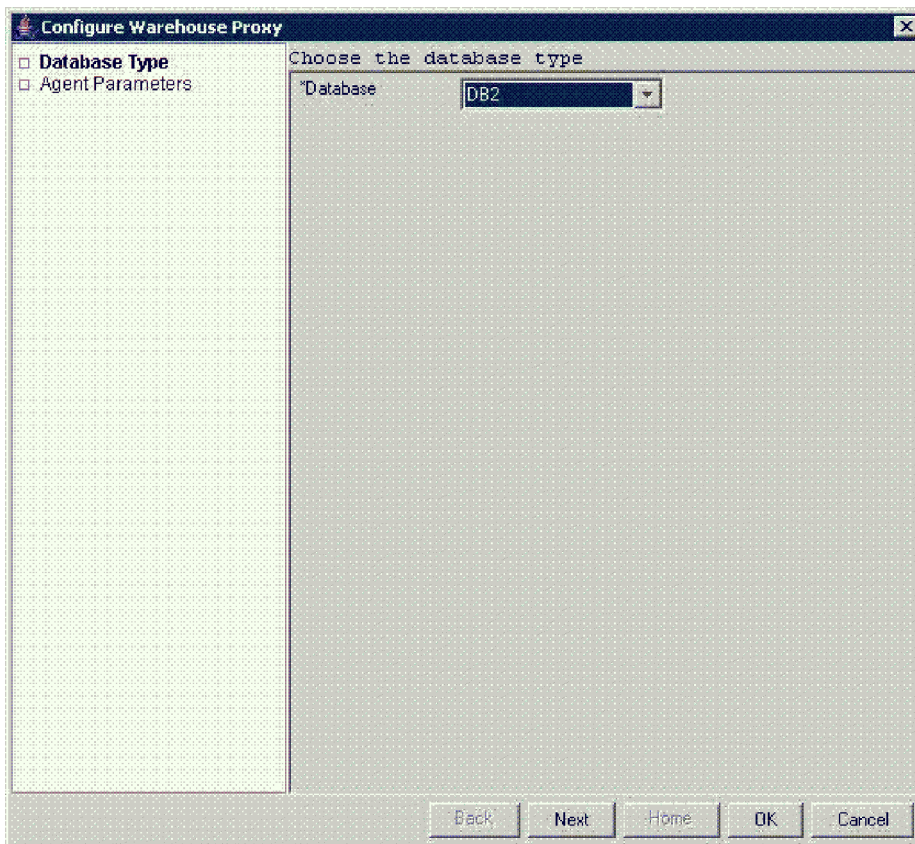


Figure 124. Warehouse Proxy Database Selection screen

The following configuration window is displayed.

Figure 125. Configure DB2 Data Source for Warehouse Proxy window

5. Click **OK** to accept all default information on this window, or change one or more default values and then click **OK**.

Note: The values for the data source name, database name, and database user ID and password must match the values that you used when configuring an ODBC connection for the Warehouse Proxy Agent. See “Configuring an ODBC data source for a DB2 data warehouse” on page 551.

Table 98. Configuration information for the Tivoli Data Warehouse database on DB2 running on z/OS

Field	Default value	Description
ODBC DSN	ITM Warehouse	The name of the data source.
Username	ITMUser	The name of the user that the Warehouse Proxy Agent will use to access the Tivoli Data Warehouse database.
Password	itmpswd1	The password for the user that the Warehouse Proxy Agent will use to access the Tivoli Data Warehouse database. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Confirm Password	itmpswd1	Confirm the password by entering it again.
Test Connection		Test the connection to the Tivoli Data Warehouse Database based on the completed fields above: ODBC DSN, Username, and Password.

Table 98. Configuration information for the Tivoli Data Warehouse database on DB2 running on z/OS (continued)

Field	Default value	Description
Warehouse TEMS List		Environment variable containing a space delimited list of TEMS names, which are given during the configuration of the HTEMS or an RTEMS. A Tivoli Enterprise Monitoring Server name in this field indicates that all the agents connected to this Tivoli Enterprise Monitoring Server will have their historical data sent to this Warehouse Proxy Agent. This variable is used when the IBM Tivoli Monitoring environment contains multiple Warehouse Proxy Agents and the workload has to be balanced using specific Warehouse Proxy Agents.
Use Batch		Batch inserts increase the data insertion rate of the Warehouse Proxy Agent. This is especially true if the proxy and the warehouse are located on different hosts. Batch inserts are supported for both ODBC and JDBC warehouse connections. Using batch inserts is recommended in all configurations, but they increase the load on the data warehouse.
Database Compression		If the database compression mode is supported by the database, the Warehouse Proxy Agent creates all tables and indexes in the Tivoli Data Warehouse with compression enabled. This option reduces the storage costs of the Tivoli Data Warehouse.
Warehouse Compression for Z/OS Sources		Select this option for the Warehouse Proxy server to allow clients installed on z/OS machines to send compressed data.
Warehouse Compression for Distributed Sources		Select this option for the Warehouse Proxy server to allow clients installed on distributed machines (those running Linux/UNIX or Windows) to send compressed data.

6. Click **OK**.

On Linux/UNIX:

1. Copy the DB2 JAR files from the DB2 server or DB2 connect to the local computer where the Warehouse Proxy Agent is installed. You can copy these files to any directory on the local computer.
db2jcc.jar
db2jcc_license_cisuz.jar
2. For instructions on configuring a Warehouse Proxy Agent on Linux or UNIX, see “Configuring a Warehouse Proxy Agent on Linux or UNIX (JDBC connection)” on page 605.

A Warehouse Proxy Agent on Linux or AIX communicates with the warehouse database through a JDBC connection. Install a Type 4 driver (DB2 on z/OS JDBC Universal Driver) on the computer where the Warehouse Proxy Agent is located. The default DB2 JDBC Driver is: com.ibm.db2.jcc.DB2Driver.

Summarization and Pruning Agent (JDBC connection):

1. The JDBC driver JAR files for your database platform must be located on the computer where you installed the Summarization and Pruning Agent. The files are named db2jcc.jar and db2jcc_license_cisuz.jar. Use a Type 4 JDBC driver. Do not use the Type 2 driver.
2. Now you can configure the Summarization and Pruning Agent. You must add the names and directory locations of the JDBC driver JAR files to the JDBC Drivers list box. You then verify the default values for the database platform that are displayed in the other text fields on the Sources pane. For instructions on configuring the Summarization and Pruning Agent, see “Configuring the Summarization and Pruning Agent (JDBC connection)” on page 643.

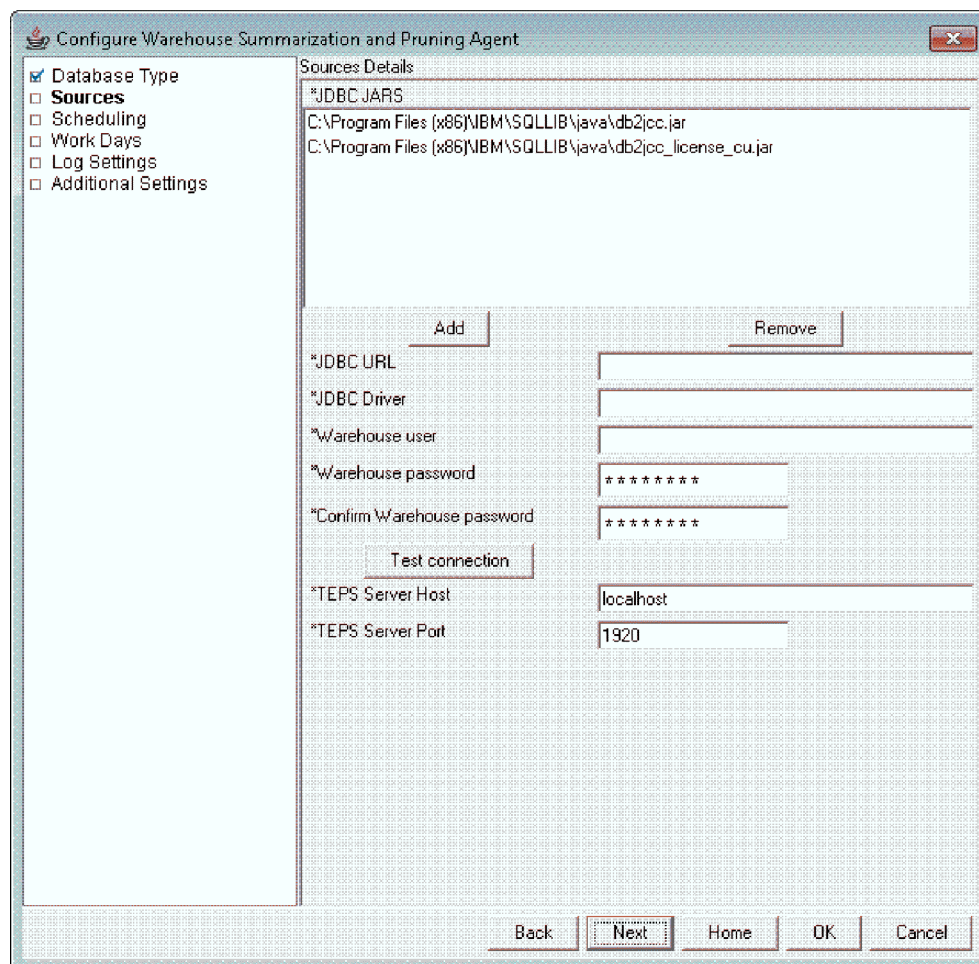


Figure 126. Sources pane of Configure Summarization and Pruning Agent window

Chapter 25. Tivoli Data Warehouse solution using Microsoft SQL Server

Use the information and instructions in this chapter to implement a Tivoli Data Warehouse solution using Microsoft SQL Server for the warehouse database.

The following table lists the goals for achieving a Microsoft SQL solution.

Table 99. Goals for achieving a Tivoli Data Warehouse solution using Microsoft SQL Server

Goal	Where to find information
Review your options, specific to a Microsoft SQL solution, for operating system platforms and communications between warehousing components.	"Supported components"
Complete prerequisite configuration steps before implementing your Tivoli Data Warehouse solution.	"Prerequisite installation" on page 597
Understand how to use the instructions for implementing your Tivoli Data Warehouse solution.	"Implementing a Tivoli Data Warehouse solution using Microsoft SQL Server" on page 598
Complete the steps for implementing your Tivoli Data Warehouse solution using Microsoft SQL Server for the data warehouse.	"Step 1: Create the Tivoli Data Warehouse database" on page 599 "Step 2: Install and configure communications for the Warehouse Proxy Agent" on page 600 "Step 3: Configure communications between the Tivoli Enterprise Portal Server and the data warehouse" on page 609 "Step 4: Install and configure communications for the Summarization and Pruning Agent" on page 612 "Step 5: Install and configure communications for Tivoli Performance Analyzer" on page 613

Supported components

This section illustrates the warehouse solution using Microsoft SQL Server.

Figure 127 on page 596 presents the options for a Tivoli Data Warehouse solution using Microsoft SQL Server for the warehouse database. The diagram summarizes the supported operating system platforms for the various warehousing components, the supported database products, and the connections between components. For more specific information about supported operating systems and database products, including product names and versions, see Chapter 3, "Hardware and software requirements," on page 75.




Figure 127. Tivoli Data Warehouse solution using Microsoft SQL Server

Note: An asterisk (*) next to a database client indicates that you must manually install the client if it does not already exist.

In the following discussion, numbered product components correspond to the numbers on the diagram.

1 Tivoli Data Warehouse on Microsoft SQL Server

A Tivoli Data Warehouse database on Microsoft SQL Server can be installed on supported Windows platforms.

2 Warehouse Proxy Agent

A Warehouse Proxy Agent on Linux or AIX communicates with the warehouse database through a JDBC connection. Install a Microsoft SQL Type 4 driver on the computer where the Warehouse Proxy is located.

Important: Use the 2005 SQL driver even if you are connecting to a warehouse database that was created in Microsoft SQL Server 2000.

A Warehouse Proxy Agent on Windows communicates with the warehouse database through an ODBC connection. The ODBC driver is included with the Microsoft SQL Server client. If the Tivoli Data Warehouse is located on a remote computer, install a Microsoft SQL Server client on the local computer where the Warehouse Proxy Agent is located. Also, configure a remote client connection to the Tivoli Data Warehouse.

3 Tivoli Enterprise Portal Server

A Tivoli Enterprise Portal Server on Windows can connect to a Microsoft SQL Server data warehouse through a Microsoft SQL Server client installed on the portal server. If the *portal server database* (designated as *TEPS database* in the diagram) uses Microsoft SQL Server, the client already exists. *Manually* install a Microsoft SQL Server client on the portal server only if the portal server database uses DB2 for Linux, UNIX and Windows.

The portal server communicates with the warehouse database through an ODBC connection. The ODBC driver is included with the Microsoft SQL Server client. Configure a remote client connection to the Tivoli Data Warehouse.

4 Summarization and Pruning Agent

The Summarization and Pruning Agent communicates with the warehouse database through a JDBC connection from any supported operating system. Install a Microsoft SQL Type 4 JDBC driver on the computer where the Summarization and Pruning Agent is located.

Important: Use the 2005 SQL driver even if you are connecting to a warehouse database that was created in Microsoft SQL Server 2000.

Prerequisite installation

Before you implement your Tivoli Data Warehouse solution, complete one or more hub installations, *excluding the warehousing components*.

Include the following components in each hub installation:

- The hub Tivoli Enterprise Monitoring Server
- *(Optional)* One or more remote monitoring servers
- The Tivoli Enterprise Portal Server, including the prerequisite RDBMS for the portal server database (DB2 for Linux, UNIX and Windows or Microsoft SQL Server)
- A Microsoft SQL Server instance on the computer where you will create the Tivoli Data Warehouse database. (The Tivoli Data Warehouse database can be shared in a multi-hub installation or dedicated to a single hub.) The SQL Server instance must be patched to the current service pack level.
- *(Optional)* A portal desktop client
- *(Optional)* Monitoring agents, and the application support for the monitoring agents

Note: The term *monitoring agent*, as used here, refers to agents that collect data directly from managed systems, not the Warehouse Proxy Agent or Summarization and Pruning Agent.

- *(Optional)* The Tivoli Performance Analyzer
- *(Optional)* Language packs for all languages other than English

See the following table for related information:

Table 100. Information topics related to installation of prerequisite software for a Tivoli Data Warehouse solution

Topic	Where to find information
Single and multiple hub installations	To understand the terminology related to single and multiple hub installations, see “Locating and sizing the hub Tivoli Enterprise Monitoring Server” on page 35.
Installation procedures for prerequisite components	The detailed instructions for installing the prerequisite components are described in Chapter 10, “Installing IBM Tivoli Monitoring,” on page 193. See your database documentations for instructions on how to install a supported database server.
Supported RDBMS versions	For specific information about the supported database platforms for the portal server database and the Tivoli Data Warehouse, see Chapter 3, “Hardware and software requirements,” on page 75.

Implementing a Tivoli Data Warehouse solution using Microsoft SQL Server

The implementation instructions in this chapter are based on certain assumptions.

Use the instructions in the remainder of this chapter to implement a Tivoli Data Warehouse solution using Microsoft SQL Server for the data warehouse.

Assumptions

The implementation instructions are based on the following assumptions:

- You will create the Tivoli Data Warehouse database on a different computer from the Tivoli Enterprise Portal Server.
- You will create a single Tivoli Data Warehouse database, to be used either within a single hub installation or to be shared in a multi-hub installation. If you have multiple independent hub installations, repeat the implementation steps for each hub installation. (See “Locating and sizing the hub Tivoli Enterprise Monitoring Server” on page 35 for information about hub installations.)
- No assumption is made about where you will install the Warehouse Proxy Agent and Summarization and Pruning Agent. Either of these agents may be installed on the same computer as the Tivoli Data Warehouse or on a different computer.

Solution steps

To implement your Tivoli Data Warehouse solution using Microsoft SQL Server, complete the four major steps described in the remaining sections of this chapter, in the order listed.

About this task

Except for Step 1, each major step consists of a series of installation and configuration tasks, listed and described in a table. Use the step tables as a road map for implementing your solution. The step tables describe the tasks at a high level, account for variations among configuration options (such as which operating system is used for a component), and reference the appropriate sections for detailed implementation procedures. To implement your solution successfully:

- Perform the tasks in the order listed in the table.
- Do not skip a table to the procedures that follow it.

Be aware that some of the implementation procedures referenced in a table are included in this chapter and some are documented elsewhere. In some cases, the task is described in the table, without referencing a separate procedure. Read and follow all instructions in the tables.

Procedure

1. Create the Tivoli Data Warehouse database.
2. Install and configure communications for the Warehouse Proxy Agent.
3. Configure communications between the Tivoli Enterprise Portal Server and the data warehouse.
4. Install and configure communications for the Summarization and Pruning Agent.

Step 1: Create the Tivoli Data Warehouse database

This section provides guidelines for creating the Tivoli Data Warehouse database using Microsoft SQL Server 2000, 2005, and 2008.

For specific instructions on how to create a Microsoft SQL Server database, see the Microsoft SQL Server documentation or have a database administrator create the database for you.

When you create the warehouse database using Microsoft SQL Server, follow these guidelines:

- Connect to the Microsoft SQL database server and create the Tivoli Data Warehouse database using the system administrator (sa) user.
- Create a database user login name and password that the warehousing components (portal server, Warehouse Proxy Agent, and Summarization and Pruning Agent) can use to access the data warehouse. In these instructions, this user account is referred to as the *warehouse user*.

You must have SQL Server authentication to create the warehouse user.

Note: The warehousing components must *not* use the system administrator (sa) user to connect to the data warehouse.

- Consider using the default values shown in the following table for the warehouse name and warehouse user. The default values are used in the configuration procedures for connecting the warehousing components to the warehouse database.

Table 101. Default values for Tivoli Data Warehouse parameters

Parameter	Default value
Tivoli Data Warehouse database name Note: When connecting to a DB2 on z/OS database, this value is unnecessary.	WAREHOUS
User name	ITMUser
User password	itmpswd1

- If the Warehouse Proxy and Summarization and Pruning Agents create database objects at runtime, you must give the warehouse user **public** and **db_owner** privileges to the Tivoli Data Warehouse database.

The warehouse user may have much fewer rights if the schema publication tool is used to create the database objects. If the schema tool is used, the warehouse user needs only the **db_datareader** and **db_datawriter** roles. If the warehouse user has limited privileges, the schema tool must be used to create any additional database objects (using the schema tool's updated mode) if the historical configuration is changed; see Chapter 20, "Schema Publication Tool," on page 525.

- For Microsoft SQL Server 2005 and 2008, do the following:
 - Create a schema with the same name (and owner) as the database user login name (for example, ITMUser) and change the default schema for the user from dbo to this login name. (This step is not necessary if you are using Microsoft SQL Server 2000.)
 - Make sure the database is set up to support inbound network TCP/IP connections.

Limiting the authority of the warehouse user

If you do not want the warehouse user to have broad administrative authority, you can limit the authority of the warehouse user to just those privileges required for accessing and using the data warehouse.

About this task

The warehouse user only needs authority for the following activities in the warehouse database:

- Create tables, alter tables, and create indexes on tables.
- Create views.
- Create functions.

To limit the authority of the warehouse user, perform the following steps in either the Microsoft SQL Server Management Studio GUI or the `osql` command-line. You must be logged into the Tivoli Data Warehouse database as a user with administrative privileges:

Procedure

1. Remove all fixed database and server roles from the warehouse user such as `db_datareader` and `db_datawriter`. If using the command-line, use the `sp_helpuser` stored procedure to see what roles are assigned to the warehouse user and `sp_droprolemember` and `sp_dropsrvrolemember` to remove the warehouse user from roles.
2. Grant the minimal necessary privileges to the warehouse user. Issue the SQL command: `grant create table, create function, create view to <warehouse user>`. Where `<warehouse user>` is the warehouse user name.

What to do next

Perform these additional steps if your security policy prohibits use of the create table, create view and create function privileges:

1. Use the schema tool to generate the DDL that create the database objects.

Note: You should create the historical collections that you want first and then configure the Summarization and Pruning Agent so that you can use the schema tool's configured mode.

2. Execute the generated scripts as described in Chapter 20, "Schema Publication Tool," on page 525. Run the scripts as the Tivoli Data Warehouse user so that the Tivoli Data Warehouse user has sufficient privileges on the tables.
3. Revoke the create table, create view, and create function privileges from the Tivoli Data Warehouse user.

Step 2: Install and configure communications for the Warehouse Proxy Agent

You can install one or more Warehouse Proxy Agents to collect and send historical data to the Tivoli Data Warehouse database.

Complete the tasks described in the following table, in the order listed, to install and configure each Warehouse Proxy Agent.

Table 102. Tasks for installing and configuring communications for the Warehouse Proxy Agent

Task	Procedure
<p>Install one or more Warehouse Proxy Agents. If you want to install a Summarization and Pruning Agent on the same computer as one of the Warehouse Proxy Agents, use the referenced procedures to install both agents at the same time.</p> <p>If you are installing more than one Warehouse Proxy Agent, each agent must be installed on a separate computer.</p> <p>The installation procedure for Windows includes steps for configuring the connection between the Warehouse Proxy and the hub Tivoli Enterprise Monitoring server. On Linux or AIX, this step is performed in a separate configuration procedure (<i>Configuring the monitoring agent</i>). See the information at right. Be sure to perform all of the referenced installation and configuration procedures.</p> <p>Note for sites setting up autonomous operation: The installation procedure includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring Server. On Windows operating systems, if you want to run the Warehouse Proxy Agent without a connection to the hub, accept the defaults for the connection information, but specify a non valid name for the monitoring server. On UNIX and Linux operating systems, check No TEMS on the TEMS Connection tab of the configuration window.</p>	<p>To install a Warehouse Proxy Agent on Windows, complete the procedure “Windows: Installing a monitoring agent” on page 262.</p> <p>To install a Warehouse Proxy Agent on Linux or AIX, complete the procedure “Linux or UNIX: Installing a monitoring agent” on page 267, including the following subsections:</p> <ul style="list-style-type: none"> • <i>Installing the monitoring agent</i> • <i>Configuring the monitoring agent</i> • <i>Changing the file permissions for agents</i> (if you used a non-root user to install the Warehouse Proxy) <p><i>Do not complete the procedure for starting the agent.</i></p>
<p>(Warehouse Proxy Agent on Windows only)</p> <ul style="list-style-type: none"> • Install a Microsoft SQL Server client on the computer where the Warehouse Proxy Agent is installed if <i>both</i> of the following statements are true: <ul style="list-style-type: none"> – The Warehouse Proxy is installed on Windows, and – The Warehouse Proxy needs to connect to a remote data warehouse. • Configure a remote client connection to the data warehouse server using Microsoft SQL Server tools. 	<p>See the Microsoft SQL Server documentation for instructions on how to install a Microsoft SQL client and configure a remote client connection.</p>
<p>(Warehouse Proxy Agent on Windows only)</p> <p>On the computer where the Warehouse Proxy Agent is installed, configure an ODBC data source for the data warehouse.</p> <p>Perform this procedure whether or not the Warehouse Proxy Agent and the warehouse database are installed on the same computer.</p>	<p>“Configuring an ODBC data source for a Microsoft SQL data warehouse” on page 602</p>
<p>(Warehouse Proxy Agent on Linux or AIX only)</p> <p>Install the most current SQL Server JDBC driver on the computer where the Warehouse Proxy Agent is installed.</p>	<p>Go to the Microsoft Web page at:</p> <p>http://www.microsoft.com and search for JDBC driver.</p> <p>Follow the instructions on the Microsoft download page for installing the driver. After you install the driver, the JAR file name and location are as follows:</p> <pre><mssqlinstalldir>/sqljdbc_1.1/enu/sqljdbc4.jar</pre>

Table 102. Tasks for installing and configuring communications for the Warehouse Proxy Agent (continued)

Task	Procedure
Configure the Warehouse Proxy Agent to connect to the data warehouse. Perform this procedure whether or not the Warehouse Proxy Agent and the warehouse database are installed on the same computer.	For a Warehouse Proxy Agent on Windows, see “Configuring a Warehouse Proxy Agent on Windows (ODBC connection)” on page 603. For a Warehouse Proxy Agent on Linux or AIX, see “Configuring a Warehouse Proxy Agent on Linux or UNIX (JDBC connection)” on page 605.
If you are installing more than one Warehouse Proxy Agent within the same hub monitoring server installation, associate each Warehouse Proxy Agent with a subset of monitoring servers (hub or remote) within the installation. Each Warehouse Proxy Agent receives data from the monitoring agents that report to the monitoring servers on the list. Use the environment variable KHD_WAREHOUSE_TEMS_LIST to specify a list of monitoring servers to associate with a Warehouse Proxy Agent.	For instructions about installing and configuring multiple Warehouse Proxy Agents within a single hub monitoring server installation, see “Installing and configuring multiple Warehouse Proxy Agents” on page 656.
(Optional) Customize the configuration of the Warehouse Proxy Agent for tuning performance.	“Tuning the performance of the Warehouse Proxy” on page 666
Start the Warehouse Proxy Agent.	“Starting the Warehouse Proxy Agent” on page 609

Configuring an ODBC data source for a Microsoft SQL data warehouse

A Microsoft SQL client on Windows requires an ODBC connection to the data warehouse. For the Warehouse Proxy Agent, you must configure the ODBC connection manually.

Before you begin

This procedure uses default values for the data source name and warehouse user ID. (Default values are used in configuration procedures for warehousing components.) Substitute different values if you do not want to use the default values.

About this task

Complete the following procedure to set up an ODBC connection for a Warehouse Proxy Agent on Windows to a local or remote Tivoli Data Warehouse.

Procedure

1. Open the Control Panel.
2. Click **Administrative Tools** → **Data Sources (ODBC)**
3. Click **Add** in the **System DSN** tab in the ODBC Data Source Administrator window.
4. Select **SQL Server** and click **Finish**.
5. Enter ITM Warehouse in the **Name** field.
6. Select the Microsoft SQL Server where the Tivoli Data Warehouse is located from the drop down list and click **Next**.
7. Select **With SQL Server authentication using a login ID and password entered by the user**.
8. Enter ITMUser for the **Login ID**. The user ID must match exactly, including case, what was created in the SQL Server database if using database authentication or the ID that was created in the Windows OS. Mixing the case will cause tables to be created with dbo as the owner rather than ITMUser. This will cause many warehousing components to not work correctly.
9. Type a password for the user in the **Password** field. The default password is itm1pswd1.

10. Click **Next**.
11. Click **Next** again.
12. Click **Finish**.
13. Click **Test Data Source** to test the connection to the database.
14. Click **OK**.

Configuring a Warehouse Proxy Agent on Windows (ODBC connection)

Use this procedure to configure a Warehouse Proxy Agent on Windows to connect to a Tivoli Data Warehouse in Microsoft SQL Server.

Procedure

1. Log on to the Windows system where the Warehouse Proxy Agent is installed and begin the configuration:
 - a. Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Monitoring Services**. The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Warehouse Proxy** and click **Configure Using Defaults**. Click **Reconfigure** if the Warehouse Proxy is installed on the same computer as the portal server.
 - c. Click **OK** on the message regarding connection to a hub monitoring server.
2. The next two windows (entitled Warehouse Proxy: Agent Advanced Configuration) contain the settings for the connection between the Warehouse Proxy Agent and the hub monitoring server. These settings were specified when the Warehouse Proxy Agent was installed. Click **OK** on each window to accept the settings.
3. Click **Yes** on the message asking if you want to configure the ODBC data source.
4. Select **SQL Server** from the list of databases and click **Next**. The following configuration window is displayed.

Figure 128. Configure SQL Data Source for Warehouse Proxy window

5. Click **OK** to accept all default information on this window, or change one or more default values and then click **OK**. The fields on this window are described in Table 103.

Note: The values for the data source name, and database user ID and password must match the values that you used when configuring an ODBC connection for the Warehouse Proxy Agent. (See “Configuring an ODBC data source for a Microsoft SQL data warehouse” on page 602.)

Table 103. Configuration information for the Tivoli Data Warehouse database on Microsoft SQL Server

Field	Default value	Description
ODBC DSN	ITM Warehouse	The name of the data source.
Username	ITMUser	The name of the user that the Warehouse Proxy Agent will use to access the Tivoli Data Warehouse database.
Password	itmpswd1	The password that the Warehouse Proxy Agent will use to access the Tivoli Data Warehouse database. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Confirm Password	itmpswd1	Confirm the password by entering it again.
Test Connection		Test the connection to the Tivoli Data Warehouse Database based on the completed fields above: ODBC DSN, Username, and Password. Note: Test Connection is not available if configuring remotely from the Tivoli Enterprise Portal.

Table 103. Configuration information for the Tivoli Data Warehouse database on Microsoft SQL Server (continued)

Field	Default value	Description
Warehouse TEMS List		Environment variable containing a space delimited list of TEMS names, which are given during the configuration of HTEMS or RTEMS. A TEMS name in this field indicates that all the agents connected to this TEMS will have their historical data sent to this Warehouse Proxy Agent. This variable is used when the ITM environment contains multiple Warehouse Proxy Agents and the workload has to be balanced using specific Warehouse Proxy Agents.
Use Batch		Batch inserts can greatly increase the data insertion rate of the Warehouse Proxy Agent. This is especially true if the proxy and the warehouse are located on different hosts. Batch inserts are supported for ODBC warehouse connections. Using batch inserts is recommended in all configurations, but they will place increased load on the data warehouse.
Database Compression		If the database compression mode is supported by the database, the Warehouse Proxy Agent will create all tables and indexes in the Tivoli Data Warehouse with compression enabled. This option reduces the storage costs of the Tivoli Data Warehouse.
Warehouse Compression for Z/OS Sources		Select this option for the Warehouse Proxy server to allow clients installed on Z/OS machines to send compressed data.
Warehouse Compression for Distributed Sources		Select this option for the Warehouse Proxy server to allow clients installed on distributed machines (Linux/UNIX, Windows) to send compressed data.

6. Click OK.

Configuring a Warehouse Proxy Agent on Linux or UNIX (JDBC connection)

Use this procedure to configure a Warehouse Proxy Agent on Linux or UNIX to connect to a Microsoft SQL Server data warehouse.

Procedure

1. Log on to the computer where the Warehouse Proxy Agent is installed and begin the configuration.
 - a. Change to the *install_dir/bin* directory and run the following command:


```
./itmcmd manage [-h install_dir]
```

where *install_dir* is the installation directory for IBM Tivoli Monitoring. The default installation directory is */opt/IBM/ITM*.

The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Warehouse Proxy** and click **Configure**. The Configure Warehouse Proxy window is displayed.

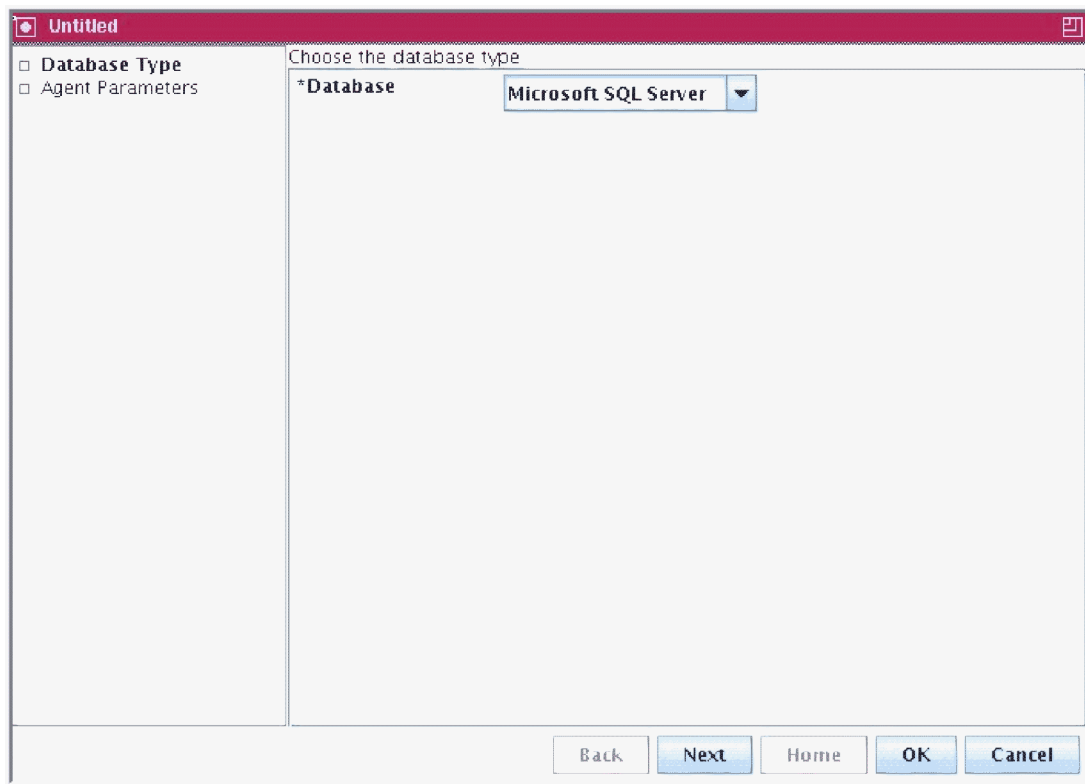


Figure 129. Configure Warehouse Proxy window (Database Type)

2. Select **Microsoft SQL Server** from the list of selectable databases, and click **Next**. The following configuration window is displayed.




Figure 130. Configure Warehouse Proxy window (Agent Parameters tab)

3. Review the settings for the connection between the Warehouse Proxy Agent and the hub monitoring server. Correct the settings if necessary. The Warehouse Proxy Agent must use the same protocols used by the application agents and by the hub monitoring. If the proxy agent does not have the same protocol as the hub monitoring server, it cannot register with the hub. If the proxy does not have the same protocol as the application agents, then the application agents cannot communicate with the proxy when they to create a route to it.
4. Add the name and directory location of the JDBC driver JAR file to the **JDBC Drivers** list box:
 - a. Click **Add** to display the file browser window. Navigate to the location of the JDBC driver JAR file on this computer and select the sqljdbc4.jar file.

Important: Use the latest SQL Server JDBC driver from Microsoft, which supports SQL Server 2008, 2005 and 2000.

- b. Click **OK** to close the browser window and add the driver file to the list.

If you need to delete an entry from the list, select the entry and click **Remove**.

5. Change the default value displayed in the **Warehouse URL** field if it is not correct. The Warehouse URL for Microsoft SQL Server 2005 is as follows:

```
jdbc:sqlserver://localhost:1433;databaseName=WAREHOUS;SelectMethod=cursor
```

- If the Tivoli Data Warehouse is installed on a remote computer, specify the host name of the remote computer instead of localhost.
- Change the port number if it is different.
- If the name of the Tivoli Data Warehouse database is not WAREHOUS, replace WAREHOUS with the actual name. (See Table 101 on page 599.)

6. Verify the JDBC driver name, which is displayed in the **JDBC Driver** field. (Note that the **JDBC Driver** field displays the *driver name*, in contrast to the *JDBC JARS* that are listed in the **JDBC JARS** field.) The Microsoft SQL Server Driver name is:

com.microsoft.sqlserver.jdbc.SQLServerDriver

7. If necessary, change the entries in the **Username** and **Password** fields to match the user name and password that were created for the Tivoli Data Warehouse. (See “Step 1: Create the Tivoli Data Warehouse database” on page 599.) The default user name is ITMUser and the default password is itm1.
8. Click **Test connection** to ensure you can communicate with the Tivoli Data Warehouse database.
9. A Tivoli Enterprise Monitoring Server name in the **Warehouse TEMS List** field indicates that all the agents connected to this Tivoli Enterprise Monitoring Server will have its historical data sent to this Warehouse Proxy Agent. This variable is used when the IBM Tivoli Monitoring environment contains multiple Warehouse Proxy Agents and the workload has to be balanced using specific Warehouse Proxy Agents.
10. Select the **Use Batch** check box if you want the Warehouse Proxy Agent to submit multiple execute statements to the Tivoli Data Warehouse database for processing as a batch. In some situations, such as crossing a network, sending multiple statements as a unit is more efficient than sending each statement separately. Batch processing is one of the features provided with the JDBC 2.0 API.
11. Select the **Database Compression** check box for the Warehouse Proxy Agent to create all tables and indexes in the Tivoli Data Warehouse with compression enabled, if the database compression mode is supported by the database. This option reduces the storage costs of the Tivoli Data Warehouse.
12. Select the **Warehouse Compression for Z/OS Sources** check box for the Warehouse Proxy server to allow clients installed on Z/OS machines to send compressed data.
13. Select the **Warehouse Compression for Distributed Sources** check box for the Warehouse Proxy server to allow clients installed on distributed machines to send compressed data.
14. Click **Save** to save your settings and close the window.

Configuring the Warehouse Proxy Agent on Linux or UNIX: command-line procedure

Complete the steps in this section to configure the Warehouse Proxy Agent from the command-line on Linux or UNIX.

Procedure

1. Log on to the computer where the Warehouse Proxy Agent is installed.
2. At the command-line change to the *ITMinstall_dir/bin* directory, where *ITMinstall_dir* is the directory where you installed the product.
3. Run the following command to start configuring the Warehouse Proxy Agent:

```
./itmcmd config -A hd
```

where *hd* is the product code for the Warehouse Proxy Agent.

Results

Here is a sample of Warehouse Proxy Agent configuration from the command-line:

```
itmcmd config -A hd
```

Database Type

Database [1=DB2, 2=Oracle, 3=Microsoft SQL Server] (default is: 1):

Agent Parameters :

Fully qualified paths to JDBC JAR files (comma separated)

JDBC JARS List (default is:): /data/jdbc/db2jcc.jar,/data/jdbc/db2jcc_license_cu.jar

The Warehouse JDBC URL
 JDBC URL (default is: jdbc:db2://localhost:50000/WAREHOUSE):

The Warehouse JDBC Driver
 JDBC Driver (default is: com.ibm.db2.jcc.DB2Driver):

The Warehouse database username
 Username (default is: ITMUSER):

The Warehouse database user password
 Enter Password (default is:):
 Re-type : Password (default is:):

Space or comma separated list of Tivoli Enterprise Monitoring Server instances served by this Warehouse Proxy agent.
 *ANY can be specified if this Warehouse Proxy agent will export data of any agents connected to any TEMS. If the list is left blank, this Warehouse Proxy agent will be the default Warehouse proxy agent.
 Warehouse TEMS List (default is:): REMOTE_ITMTDWP12

Batch Database Operations
 Use Batch [1=TRUE, 2=FALSE] (default is: 1):

Database Compression option
 Database Compression [1=TRUE, 2=FALSE] (default is: 2):

Enable the compression of historical data from Z/OS sources before upload to the Warehouse Proxy Server
 Warehouse Compression for Z/OS Sources [1=TRUE, 2=FALSE] (default is: 2):

Enable the compression of historical data from distributed sources before upload to the Warehouse Proxy Server
 Warehouse Compression for Distributed Sources [1=TRUE, 2=FALSE] (default is: 1):

Will this agent connect to a TEMS? [1=YES, 2=NO] (Default is: 1): 1
 TEMS Host Name (Default is: itmtdwp18):

Starting the Warehouse Proxy Agent

This section contains instructions for starting the Warehouse Proxy Agent.

- To start the Warehouse Proxy Agent from the Manage Tivoli Enterprise Monitoring Services window, right-click **Warehouse Proxy** and select **Start**.
- (*Linux or AIX only*) To start the Warehouse Proxy Agent from the command-line, run the following command from the bin directory of the IBM Tivoli Monitoring installation directory. The default installation directory is /opt/IBM/ITM.
 ./itmcmd agent start hd

where hd is the product code for the Warehouse Proxy Agent.

Step 3: Configure communications between the Tivoli Enterprise Portal Server and the data warehouse

Complete the tasks in this section, in the order listed, to configure communications between the portal server and the data warehouse.

Table 104. Tasks for configuring communications between the portal server and a Microsoft SQL Server data warehouse

Task	Procedure
If the <i>portal server database</i> was created using DB2 for Linux, UNIX and Windows, install a Microsoft SQL Server database client on the portal server.	See the Microsoft SQL Server documentation for instructions on how to install a Microsoft SQL Server database client.
If the portal server database was created using Microsoft SQL Server, the Microsoft SQL Server database client already exists on the portal server.	
Configure a remote client connection to the data warehouse server using Microsoft SQL Server tools.	See the Microsoft SQL Server documentation for instructions on how to configure a remote client connection.

Table 104. Tasks for configuring communications between the portal server and a Microsoft SQL Server data warehouse (continued)

Task	Procedure
Configure the portal server to connect to the data warehouse. The configuration procedure automatically configures an ODBC connection to the data warehouse.	"Configuring the portal server (ODBC connection)"
Restart the portal server.	On the Manage Tivoli Enterprise Monitoring Services window, right-click Tivoli Enterprise Portal Server and select Start .
Test the connection between the portal server and the Tivoli Data Warehouse by creating a customized query in the Tivoli Enterprise Portal.	"Testing the connection between the portal server and the Tivoli Data Warehouse" on page 663

Configuring the portal server (ODBC connection)

Use the procedure in this section to configure the portal server to connect to a Tivoli Data Warehouse in Microsoft SQL Server.

Procedure

- Log on to the Windows system where the portal server is installed and begin the configuration:
 - Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Monitoring Services**. The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - Right-click Tivoli Enterprise Portal Server and click **Reconfigure**.
- The next two windows (entitled TEP Server Configuration) contain the settings for the connection between the portal server and the hub monitoring server. These settings were specified when the portal server was installed. Click **OK** on each window to accept the settings.
- Click **Yes** on the message asking if you want to reconfigure the warehouse information for the Tivoli Enterprise Portal Server.
- Select **SQL Server** from the list of databases and click **OK**. The following configuration window is displayed.

Configure Warehouse Data Source for TEPS - SQL Server

Data Source Name:

Please enter the Database User ID and Password required for connecting to the Warehouse Data Source:

Database User ID:

Database Password:

Reenter Password:

☐ The database and the Warehouse Data Source do not exist and should be created by the Installer. Note: the SQL Server must be installed on this machine.

Database Name:

Admin User ID:

Admin Password:

OK Cancel Help

Figure 131. Configure SQL Data Source for Warehouse window

- Click **OK** to accept all default information on this window, or change one or more default values and then click **OK**. The fields on this window are described in the following table:

Table 105. Configuration information for the Tivoli Data Warehouse database on Microsoft SQL Server

Field	Default value	Description
Data Source Name	ITM Warehouse	The name of the data source.
Database User ID	ITMUser	The login name of the database user that the portal server will use to access the Tivoli Data Warehouse database.
Database Password	(no default)	The password for the database login user. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Reenter Password	(no default)	Confirm the password by entering it again.
Database Name	WAREHOUS	The name of the database.
Admin User ID	sa	The database administrator ID.
Admin Password	(no default)	The password for the database administrator.

- Click **OK**.

Step 4: Install and configure communications for the Summarization and Pruning Agent

Complete the tasks described in this section, in the order listed, to install and configure the Summarization and Pruning Agent.

Table 106. Tasks for installing and configuring communications for the Summarization and Pruning Agent

Task	Procedure
<p>Install the Summarization and Pruning Agent if you have not already installed it. For best performance, install the Summarization and Pruning Agent on the same computer as the data warehouse.</p> <p>The installation procedure for Windows includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring server. On Linux or AIX, this step is performed in a separate configuration procedure (<i>Configuring the monitoring agent</i>). See the information at right. Be sure to perform all referenced installation and configuration procedures.</p> <p>Note: The Summarization and Pruning Agent is not automatically started after installation. Do not complete any step or procedure for starting the agent at this point.</p>	<p>To install a Summarization and Pruning Agent on Windows, complete the procedure “Windows: Installing a monitoring agent” on page 262.</p> <p>To install a Summarization and Pruning Agent on Linux or UNIX, complete the procedure “Linux or UNIX: Installing a monitoring agent” on page 267, including the following subsections:</p> <ul style="list-style-type: none"> • <i>Installing the monitoring agent</i> • <i>Configuring the monitoring agent</i> • <i>Changing the file permissions for agents</i> (if you used a non-root user to install the Warehouse Proxy) <p><i>Do not complete the procedure for starting the agent.</i></p>
<p>Install the most current SQL Server JDBC driver on the computer where the Summarization and Pruning Agent is installed.</p>	<p>Go to the Microsoft Web page at:</p> <p>http://www.microsoft.com and search for JDBC driver.</p> <p>Follow the instructions on the Microsoft download page for installing the driver. After you install the driver, the JAR file name and location are as follows:</p> <pre><mssqlinstalldir>/sqljdbc_1.1/enu/sqljdbc4.jar</pre>
<p>Configure the Summarization and Pruning Agent.</p> <p>When you configure the Summarization and Pruning Agent, you configure the connection to the Tivoli Data Warehouse and you specify settings that control the operation of the Summarization and Pruning Agent.</p> <p>Perform this procedure whether or not the Summarization and Pruning Agent and the warehouse database are installed on the same computer.</p>	<p>“Configuring the Summarization and Pruning Agent (JDBC connection)” on page 643</p>
<p>Configure the Summarization and Pruning Agent to connect to the Tivoli Enterprise Portal Server. Perform this procedure whether or not the Summarization and Pruning Agent and the warehouse database are installed on the same computer.</p>	<p>See Step 9 of “Configuring the Summarization and Pruning Agent (JDBC connection)” on page 643.</p>
<p>Configure history collection.</p> <p>When you configure history collection, you specify settings for how often to collect, aggregate, and prune data for individual monitoring agents and attribute groups. Configure history collection from the Tivoli Enterprise Portal.</p>	<p>See the <i>IBM Tivoli Monitoring Administrator's Guide</i> for instructions on how to configure history collection.</p>

Table 106. Tasks for installing and configuring communications for the Summarization and Pruning Agent (continued)

Task	Procedure
Start the Summarization and Pruning Agent.	"Starting the Summarization and Pruning Agent" on page 656

Step 5: Install and configure communications for Tivoli Performance Analyzer

You can install Tivoli Performance Analyzer to a server which also has other Tivoli Monitoring components installed, or you can install it to a separate machine.

The installation procedure is similar to that for monitoring agents. Complete the tasks described in the following table, in the order listed, to install and configure Tivoli Performance Analyzer.

Table 107. Tasks for installing and configuring communications for the Tivoli Performance Analyzer

Task	Procedure
<p>Install Tivoli Performance Analyzer.</p> <p>The installation procedure for Windows includes steps for configuring the connection between the Tivoli Performance Analyzer and the hub Tivoli Enterprise Monitoring server. On Linux or AIX, this step is performed in a separate configuration procedure (Configuring the monitoring agent). See the information opposite. Be sure to perform all of the referenced installation and configuration procedures.</p> <p>Note: for sites setting up autonomous operation, the installation procedure includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring Server. On Windows operating systems, if you want to run Tivoli Performance Analyzer without a connection to the hub, accept the defaults for the connection information, but specify a non valid name for the monitoring server. On UNIX and Linux operating systems, select No TEMS on the TEMS Connection in the configuration window.</p>	<p>To install Tivoli Performance Analyzer on Windows, complete the procedure "Windows: Installing a monitoring agent" on page 262. To install Tivoli Performance Analyzer agent on Linux or AIX, complete the procedure "Linux or UNIX: Installing a monitoring agent" on page 267, including the following subsections:</p> <ul style="list-style-type: none"> • Installing the monitoring agent • Configuring the monitoring agent • Changing the file permissions for agents <p>Do not complete the procedure for starting the agent.</p>
Install a Microsoft SQL client and configure a remote client connection.	See the Microsoft SQL Server documentation for instructions on how to install a Microsoft SQL client and configure a remote client connection.
On the computer where Tivoli Performance Analyzer is installed, configure an ODBC data source for the data warehouse.	"Configuring an ODBC data source for a Microsoft SQL data warehouse" on page 614
Install the most current SQL Server JDBC driver on the computer where Tivoli Performance Analyzer is installed.	<p>Go to the Microsoft Web page at:</p> <p>http://www.microsoft.com and search for JDBC driver.</p> <p>Follow the instructions on the Microsoft download page for installing the driver. After you install the driver, the JAR file name and location are as follows:</p> <pre><mssqlinstalldir>/sqljdbc_1.1/enu/sqljdbc4.jar</pre>

Table 107. Tasks for installing and configuring communications for the Tivoli Performance Analyzer (continued)

Task	Procedure
Configure Tivoli Performance Analyzer to connect to the data warehouse.	For Tivoli Performance Analyzer on Windows, see “Configuring Tivoli Performance Analyzer on Windows (ODBC connection)” on page 615. For a Tivoli Performance Analyzer on Linux or AIX, see “Configuring Tivoli Performance Analyzer on Linux or UNIX (JDBC connection)” on page 616.
Start Tivoli Performance Analyzer.	“Starting Tivoli Performance Analyzer” on page 617

Configuring an ODBC data source for a Microsoft SQL data warehouse

A Microsoft SQL client on Windows requires an ODBC connection to the data warehouse. For the Tivoli Performance Analyzer, you must configure the ODBC connection manually.

Before you begin

This procedure uses default values for the data source name and warehouse user ID. (Default values are used in configuration procedures for warehousing components.) Substitute different values if you do not want to use the default values.

About this task

Complete the following procedure to set up an ODBC connection for a Tivoli Performance Analyzer on Windows to a local or remote Tivoli Data Warehouse.

Procedure

1. Open the Control Panel.
2. Click **Administrative Tools** → **Data Sources (ODBC)**
3. Click **Add** in the **System DSN** tab in the ODBC Data Source Administrator window.
4. Select **SQL Server** and click **Finish**.
5. Enter ITM Warehouse in the **Name** field.
6. Select the Microsoft SQL Server where the Tivoli Data Warehouse is located from the drop down list and click **Next**.
7. Select **With SQL Server authentication using a login ID and password entered by the user**.
8. Enter ITMUser for the **Login ID**. The user ID must match exactly, including case, what was created in the SQL Server database if using database authentication or the ID that was created in the Windows OS. Mixing the case will cause tables to be created with dbo as the owner rather than ITMUser. This will cause many warehousing components to not work correctly.
9. Type a password for the user in the **Password** field. The default password is itm1pswd1.
10. Click **Next**.
11. Click **Next** again.
12. Click **Finish**.
13. Click **Test Data Source** to test the connection to the database.
14. Click **OK**.

Configuring Tivoli Performance Analyzer on Windows (ODBC connection)

Use this procedure to configure Tivoli Performance Analyzer on Windows to connect to a Tivoli Data Warehouse on Microsoft SQL Server.

Procedure

1. Log on to the Windows system where Tivoli Performance Analyzer is installed and begin the configuration:
 - a. Click Start → Programs → IBM Tivoli Monitoring → Manage Tivoli Monitoring Services. The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click Performance Analyzer and click **Configure Using Defaults**.
 - c. Click **OK** on the message regarding connection to a hub monitoring server.
2. The next two windows (entitled Performance Analyzer: Agent Advanced Configuration) contain the settings for the connection between Tivoli Performance Analyzer and the hub monitoring server. These settings were specified when Tivoli Performance Analyzer was installed. Click **OK** on each window to accept the settings.
3. Click **Yes** on the message asking if you want to configure the ODBC data source.
4. Select **ODBC** from the list of selectable agent database connection types.
5. Set your Database Type to **MSSQL**.
6. Specify the Data Source Name - Agent ODBC DSN (ITM Warehouse is the default).

Note: Tivoli Performance Analyzer does not create this DSN - it must already exist. If you are installing the agent on the same machine where TEP Server is installed, you can use the existing data source created by Tivoli Monitoring. Otherwise, you must create a new System DSN manually, prior to re-configuring Tivoli Performance Analyzer.

On 64-bit versions of Windows, data sources created by the default ODBC Data Source Administrator applet available from the Control Panel are not available for 32-bit applications. Therefore you must use the 32-bit version of the ODBC Data Source Administrator applet from <WINDOWS>\SysWOW64\odbcad32.exe.

Values for the data source name, database name, and database user ID and password must match the values that you used when configuring an ODBC connection for Tivoli Performance Analyzer. For more information, see “Configuring an ODBC data source for a DB2 data warehouse” on page 564.

7. Type the Username and Password. The entries in these fields are used to connect to the Tivoli Data Warehouse and are the same credentials as those used by the Tivoli Enterprise Portal Server, the Warehouse Proxy Agent and the Summarization and Pruning Agent to communicate with Tivoli Data Warehouse.
8. Click **Next** to proceed to the Advanced Configuration window.
9. You can enable Advanced Configuration to specify TDW Schema and The TDW database schema. If you do not select **Enable advanced configuration** these options are greyed out.
10. You can also choose whether you want the agent to Initialize PA tables and OS domain tasks.

Note: Setting Initialize PA tables to YES will remove and recreate all previously created tables deleting all user tasks and reverting each OS task to its default.

11. Use the **Bypass connection tests** option to finish the configuration without running connection tests.
12. Click **OK** to finish the configuration process.

Note: The values for the data source name, and database user ID and password must match the values that you used when configuring an ODBC connection for Tivoli Performance Analyzer. (See “Configuring an ODBC data source for a Microsoft SQL data warehouse” on page 614).

Table 108. Configuration information for the Tivoli Data Warehouse database on Microsoft SQL

Field	Default value	Description
ODBC DSN	ITM Warehouse	The name of the data source.
Username	ITMUser	The name of the Windows OS user that the Tivoli Performance Analyzer will use to access the Tivoli Data Warehouse database.
Password	itmpswd1	The password for the Windows OS user. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Test Connection		Test the connection to the Tivoli Data Warehouse Database based on the completed fields above: ODBC DSN, Username, and Password.

Configuring Tivoli Performance Analyzer on Linux or UNIX (JDBC connection)

Use this procedure to configure Tivoli Performance Analyzer on Linux or UNIX to connect to a Microsoft SQL Server data warehouse.

Procedure

- To begin the configuration, log on to the computer where Tivoli Performance Analyzer is installed.
 - Change to the `install_dir/bin` directory and run the following command:

```
./itmcmd manage [-h install_dir]
```

where `install_dir` is the installation directory for IBM Tivoli Monitoring. The default installation directory is `/opt/IBM/ITM`. The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - Right-click **Performance Analyzer** and click **Reconfigure**. The Configure Tivoli Performance Analyzer window is displayed.
- Set the Database Type to MSSQL.
- Type the username and the password. The entries in these fields are used to connect to the Tivoli Data Warehouse.
- Review all the defaults in the Agent Configuration window and change as required.
 - If the Tivoli Data Warehouse is installed on a remote computer, specify the host name of the remote computer instead of localhost.
 - Change the port number if necessary (the default port number for MSSQL is 1433).
 - If the name of the Tivoli Data Warehouse database is not WAREHOUS, replace WAREHOUS with the actual name. (See “Step 1: Create the Tivoli Data Warehouse database” on page 599.)
- Specify the JDBC Driver. The default driver name is `com.microsoft.sqlserver.jdbc.SQLServerDriver`.
- Specify the JDBC Driver Path, which should be provided as a list of JAR files with the full path separated by ":".

Note: The Microsoft SQL Server 2005 driver JAR file name and default location after downloading from the Web are as follows:

```
mssql2005install_dir/sqljdbc_1.1/enu/sqljdbc4.jar
```

Fast path: You can use the Browse button to specify the path. In such a case a file list is added at the end of the JDBC Driver Path text field, separated from the existing content by a path separator.

7. You can use the **Test connection** button to check whether the connection can be initiated.
8. Click **Next** to proceed to the Advanced Configuration window.
 - a. You can enable Advanced Configuration to specify TDW Schema and Configuration schema. If you do not select Enable advanced configuration, all of these options are greyed out.
 - b. You can also choose whether you want the agent to initialize PA tables.

Note: Setting Initialize PA tables to YES will remove and recreate all previously created tables deleting all user tasks and reverting each OS task to its default.

- c. Use the **Bypass connection tests** option to finish the configuration without running connection tests.
9. Click **Save** to save your settings and close the window.

Starting Tivoli Performance Analyzer

To start Tivoli Performance Analyzer from the Manage Tivoli Enterprise Monitoring Services window, right-click Tivoli Performance Analyzer and select Start.

To start the Tivoli Performance Analyzer agent from the command-line, run the following command from the bin directory of the IBM Tivoli Monitoring installation directory. The default installation directory is /opt/IBM/ITM.

```
./itmcmd agent start pa
```

where pa is the product code for Tivoli Performance Analyzer agent.

Chapter 26. Tivoli Data Warehouse solution using Oracle

Use the information and instructions in this chapter to implement a Tivoli Data Warehouse solution using Oracle for the warehouse database.

The following table lists the goals for achieving an Oracle solution.

Table 109. Goals for achieving a Tivoli Data Warehouse solution using Oracle

Goal	Where to find information
Review your options, specific to an Oracle solution, for operating system platforms and communications between warehousing components.	"Supported components"
Install prerequisite software before implementing your Tivoli Data Warehouse solution.	"Prerequisite installation" on page 621
Understand how to use the instructions for implementing your Tivoli Data Warehouse solution.	"Implementing a Tivoli Data Warehouse solution using Oracle" on page 622
Complete the steps for implementing your Tivoli Data Warehouse solution using Oracle for the data warehouse.	"Step 1: Create the Tivoli Data Warehouse database" on page 623 "Step 2: Install and configure communications for the Warehouse Proxy Agent" on page 624 "Step 3: Configure communications between the Tivoli Enterprise Portal Server and the data warehouse" on page 632 "Step 4: Install and configure communications for the Summarization and Pruning Agent" on page 636 "Step 5: Install and configure communications for Tivoli Performance Analyzer" on page 637

Supported components

This section illustrates the warehouse solution using Oracle for the warehouse database.

Figure 132 on page 620 presents the options for a Tivoli Data Warehouse solution using Oracle for the warehouse database. The diagram summarizes the supported operating system platforms for the various warehousing components, the supported database products, and the connections between components. For more specific information about supported operating systems and database products, including product names and versions, see Chapter 3, "Hardware and software requirements," on page 75.




Figure 132. Tivoli Data Warehouse solution using Oracle

Note: An asterisk (*) next to a database client indicates that you must manually install the client if it does not already exist.

In the following discussion, numbered product components correspond to the numbers on the diagram.

1 Tivoli Data Warehouse on Oracle

A Tivoli Data Warehouse database on Oracle can be installed on supported Windows, Linux, or any UNIX platform support by Oracle. Ensure that the Oracle listener is active in order to accept connections from an Oracle client or JDBC driver.

2 Warehouse Proxy Agent

A Warehouse Proxy Agent on Linux or AIX communicates with the warehouse database through a JDBC connection. Install an Oracle Type 4 JDBC driver on the computer where the Warehouse Proxy Agent is

located. If the Tivoli Data Warehouse is located on a remote computer, create a TNS (Transparent Network Substrate) Service Name on the local computer where the Warehouse Proxy Agent is located.

A Warehouse Proxy Agent on Windows communicates with the warehouse database through an ODBC connection. The ODBC driver is included with the Oracle client. If the Tivoli Data Warehouse is located on a remote computer, install an Oracle client on the local computer where the Warehouse Proxy Agent is located. Also, create a TNS Service Name on the local computer.

3 Tivoli Enterprise Portal Server

A Tivoli Enterprise Portal Server on Windows communicates with an Oracle data warehouse through an Oracle database client using an ODBC connection. You must manually install the Oracle client on the portal server. The ODBC driver is included with the Oracle client.

A portal server on Linux or AIX communicates with the warehouse database through a JDBC connection. Install an Oracle Type 4 JDBC driver on the portal server.

4 Summarization and Pruning Agent

The Summarization and Pruning Agent communicates with the warehouse database through a JDBC connection from any supported operating system. Install an Oracle Type 4 JDBC driver on the computer where the Summarization and Pruning Agent is located.

Prerequisite installation

Before you implement your Tivoli Data Warehouse solution, complete one or more hub installations, *excluding the warehousing components*.

Include the following components in each hub installation:

- The hub Tivoli Enterprise Monitoring Server
- *(Optional)* One or more remote monitoring servers
- The Tivoli Enterprise Portal Server, including the prerequisite RDBMS for the portal server database (DB2 for Linux, UNIX, and Windows or Microsoft SQL Server)
- An Oracle database server on the computer where you will create the Tivoli Data Warehouse database. (The Tivoli Data Warehouse database can be shared in a multi-hub installation or dedicated to a single hub.)
- *(Optional)* A portal desktop client
- *(Optional)* Monitoring agents, and the application support for the monitoring agents

Note: The term *monitoring agent*, as used here, refers to agents that collect data directly from managed systems, not the Warehouse Proxy Agent or Summarization and Pruning Agent.

- *(Optional)* The Tivoli Performance Analyzer
- *(Optional)* Language packs for all languages other than English

See the following table for related information:

Table 110. Information topics related to installation of prerequisite software for a Tivoli Data Warehouse solution

Topic	Where to find information
Single and multiple hub installations	To understand the terminology related to single and multiple hub installations, see “Locating and sizing the hub Tivoli Enterprise Monitoring Server” on page 35.

Table 110. Information topics related to installation of prerequisite software for a Tivoli Data Warehouse solution (continued)

Topic	Where to find information
Installation procedures for prerequisite components	The detailed instructions for installing the prerequisite components are described in Chapter 10, “Installing IBM Tivoli Monitoring,” on page 193. See your database documentations for instructions on how to install a supported database server.
Supported RDBMS versions	For specific information about the supported database platforms for the portal server database and the Tivoli Data Warehouse, see Chapter 3, “Hardware and software requirements,” on page 75.

Implementing a Tivoli Data Warehouse solution using Oracle

The implementation instructions in this chapter are based on certain assumptions.

Use the instructions in the remainder of this chapter to implement a Tivoli Data Warehouse solution using Oracle for the data warehouse.

Assumptions

The implementation instructions are based on the following assumptions:

- You will create the Tivoli Data Warehouse database on a different computer from the Tivoli Enterprise Portal Server.
- You will create a single Tivoli Data Warehouse database, to be used either within a single hub installation or to be shared in a multi-hub installation. If you have multiple independent hub installations, repeat the implementation steps for each hub installation. (See “Locating and sizing the hub Tivoli Enterprise Monitoring Server” on page 35 for information about hub installations.)
- No assumption is made about where you will install the Warehouse Proxy Agent and Summarization and Pruning Agent. Either of these agents may be installed on the same computer as the Tivoli Data Warehouse or on a different computer.

Solution steps

To implement your Tivoli Data Warehouse solution using Oracle, complete the four major steps described in the remaining sections of this chapter, in the order listed.

About this task

Each major step consists of a series of installation and configuration tasks, listed and described in a table. Use the step tables as a road map for implementing your solution. The step tables describe the tasks at a high level, account for variations among configuration options (such as which operating system is used for a component), and reference the appropriate sections for detailed implementation procedures. To implement your solution successfully:

- Perform the tasks in the order listed in the table.
- Do not skip a table to the procedures that follow it.

Be aware that some of the implementation procedures referenced in a table are included in this chapter and some are documented elsewhere. In some cases, the task is described in the table, without referencing a separate procedure. Read and follow all instructions in the tables.

Procedure

1. Create the Tivoli Data Warehouse database.
2. Install and configure communications for the Warehouse Proxy Agent.
3. Configure communications between the Tivoli Enterprise Portal Server and the data warehouse.

4. Install and configure communications for the Summarization and Pruning Agent.

Step 1: Create the Tivoli Data Warehouse database

This section provides guidelines for creating a Tivoli Data Warehouse database using Oracle.

Complete the tasks described in the following table to create a Tivoli Data Warehouse database using Oracle and to make it accessible to clients.

Table 111. Tasks for creating the Tivoli Data Warehouse database

Task	Procedure
Create the Tivoli Data Warehouse database on one of the supported Windows, Linux, or UNIX operating systems. To comply with the assumptions described in the introduction to this chapter, create the database on a different computer from the Tivoli Enterprise Portal Server.	For guidance on planning the size and disk requirements for the warehouse database, see “Planning considerations for the Tivoli Data Warehouse” on page 510. For information about creating the warehouse database using Oracle, see “Creating the warehouse database on Oracle.”
Activate the Oracle listener on the Oracle server where the Tivoli Data Warehouse is installed. To activate the Oracle listener: <ul style="list-style-type: none">• Use the Oracle Listener Service on Windows.• Use the lsnrctl start command on Linux and UNIX.	See the Oracle documentation for instructions on how to activate the Oracle listener.

Creating the warehouse database on Oracle

This section provides guidelines for creating the Tivoli Data Warehouse database on Oracle.

For specific instructions on how to create an Oracle database, see the Oracle documentation or have a database administrator create the database for you.

When you create the warehouse database using Oracle, follow these guidelines:

- Use the Unicode character set (AL32UTF8) when creating the database.
- Create a database user login name and password that the warehousing components (portal server, Warehouse Proxy Agent, and Summarization and Pruning Agent) can use to access the data warehouse. In these instructions, this user account is referred to as the *warehouse user*.
- Use the default values shown in the following table for the warehouse name and warehouse user. The default values are used in the configuration procedures for connecting the warehousing components to the warehouse database.

Table 112. Default values for Tivoli Data Warehouse parameters

Parameter	Default value
Tivoli Data Warehouse database name	WAREHOUS
User name	ITMUser
User password	itmpswd1

- Create an ITM_DW role, and give this role the following permissions:

```
GRANT CREATE CONNECT TO ITM_DW;  
GRANT CREATE PROCEDURE TO ITM_DW;  
GRANT CREATE TABLE TO ITM_DW;  
GRANT CREATE VIEW TO ITM_DW;
```

After you create the warehouse user ID that will be used by the Warehouse Proxy and the Summarization and Pruning Agents to connect to the Tivoli Data Warehouse database, give this user ID the *role* you just created:

```
CREATE USER itmuser PROFILE DEFAULT IDENTIFIED by itmuser_password
ACCOUNT UNLOCK;
GRANT role TO itmuser;
```

As all the Tivoli Data Warehouse tables are created in the user's default tablespace, you need to allocate enough space quota on the default tablespace to this user to create all the tables, or you can simplify it by allowing unlimited tablespace to this user.

Note: The ITM_DW role needs the connect privilege only if the database objects are created using the schema publication tool. If the historical configuration is changed and the warehouse user has limited privileges, the schema tool must be used to create any additional database objects (using the schema tool's updated mode; see Chapter 20, "Schema Publication Tool," on page 525).

- Activate the Oracle listener using the Oracle Listener Service on Windows or the **lsnrctl start** command on Linux and UNIX.

Step 2: Install and configure communications for the Warehouse Proxy Agent

You can install one or more Warehouse Proxy Agents to collect and send historical data to the Tivoli Data Warehouse database.

Complete the tasks described in the following table, in the order listed, to install and configure each Warehouse Proxy Agent.

Table 113. Tasks for installing and configuring communications for the Warehouse Proxy Agent

Task	Procedure
Install one or more Warehouse Proxy Agents. If you want to install a Summarization and Pruning Agent on the same computer as one of the Warehouse Proxy Agents, use the referenced procedures to install both agents at the same time.	To install a Warehouse Proxy Agent on Windows, complete the procedure "Windows: Installing a monitoring agent" on page 262.
If you are installing more than one Warehouse Proxy Agent, each agent must be installed on a separate computer.	To install a Warehouse Proxy Agent on Linux or AIX, complete the procedure "Linux or UNIX: Installing a monitoring agent" on page 267, including the following subsections:
The installation procedure for Windows includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring server. On Linux or AIX, this step is performed in a separate configuration procedure (<i>Configuring the monitoring agent</i>). See the information at right. Be sure to perform all referenced installation and configuration procedures.	<ul style="list-style-type: none"> • <i>Installing the monitoring agent</i> • <i>Configuring the monitoring agent</i> • <i>Changing the file permissions for agents</i> (if you used a non-root user to install the Warehouse Proxy)
Note for sites setting up autonomous operation: The installation procedure includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring Server. On Windows operating systems, if you want to run the Warehouse Proxy Agent without a connection to the hub, accept the defaults for the connection information, but specify a nonvalid name for the monitoring server. On UNIX and Linux operating systems, check No TEMS on the TEMS Connection tab of the configuration window.	<i>Do not complete the procedure for starting the agent.</i>

Table 113. Tasks for installing and configuring communications for the Warehouse Proxy Agent (continued)

Task	Procedure
<p>(Warehouse Proxy Agent on Windows only)</p> <ul style="list-style-type: none"> Install an Oracle client on the computer where the Warehouse Proxy Agent is installed if <i>both</i> of the following statements are true: <ul style="list-style-type: none"> The Warehouse Proxy is installed on Windows, and The Warehouse Proxy needs to connect to a remote data warehouse. Ensure that the latest Oracle patches are installed. Set the following system variable on the computer where the Warehouse Proxy Agent is installed. Restart the computer after setting the variable. The format of the NLS_LANG environment variable is: NLS_LANG=language_territory.charset <p>Set the language and territory to appropriate variables. For the United States this is NLS_LANG=AMERICAN_AMERICA.AL32UTF8.</p> <p>Perform the last two tasks whether or not the warehouse database is local or remote.</p>	<p>See the Oracle documentation for instructions on how to install an Oracle client.</p> <p>Obtain Oracle ODBC drivers from the following Web site:</p> <p>http://www.oracle.com/technology/software/tech/windows/odbc/htdocs/utilsoft.html</p>
<p>(Warehouse Proxy Agent on Windows only)</p> <p>If the Tivoli Data Warehouse is located on a remote computer, create a TNS (Transparent Network Substrate) Service Name on the local computer where the Warehouse Proxy Agent is located.</p>	<p>“Creating a TNS Service Name” on page 626</p>
<p>(Warehouse Proxy Agent on Windows only)</p> <p>On the computer where the Warehouse Proxy Agent is installed, configure an ODBC data source for the data warehouse, using the TNS Service Name that you created in the preceding step.</p> <p>Perform this procedure whether or not the Warehouse Proxy Agent and the warehouse database are installed on the same computer.</p>	<p>“Configuring an ODBC data source for an Oracle data warehouse” on page 627</p>
<p>(Warehouse Proxy Agent on Linux or AIX only)</p> <p>Install an Oracle Type 4 JDBC driver on the computer where the Warehouse Proxy Agent is installed.</p>	<p>Obtain the Oracle JDBC Driver from the following Web site:</p> <p>http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html</p> <p>The Oracle JDBC driver JAR file name and location after installation is as follows:</p> <p><i>oracleinstall_dir/jdbc/lib/ojdbc14.jar</i></p> <p>The ojdbc14.jar file supports JRE 1.5 or higher, the required Java Runtime Environment for IBM Tivoli Monitoring.</p>
<p>Configure the Warehouse Proxy Agent to connect to the data warehouse.</p> <p>Perform this procedure whether or not the Warehouse Proxy Agent and the warehouse database are installed on the same computer.</p>	<p>For a Warehouse Proxy Agent on Windows, see “Configuring a Warehouse Proxy Agent on Windows (ODBC connection)” on page 627.</p> <p>For a Warehouse Proxy Agent on Linux or AIX, see “Configuring a Warehouse Proxy Agent on Linux or UNIX (JDBC connection)” on page 629.</p>

Table 113. Tasks for installing and configuring communications for the Warehouse Proxy Agent (continued)

Task	Procedure
If you are installing more than one Warehouse Proxy Agent within the same hub monitoring server installation, associate each Warehouse Proxy Agent with a subset of monitoring servers (hub or remote) within the installation. Each Warehouse Proxy Agent receives data from the monitoring agents that report to the monitoring servers on the list. Use the environment variable KHD_WAREHOUSE_TEMS_LIST to specify a list of monitoring servers to associate with a Warehouse Proxy Agent.	For instructions about installing and configuring multiple Warehouse Proxy Agents within a single hub monitoring server installation, see “Installing and configuring multiple Warehouse Proxy Agents” on page 656.
(Optional) Customize the configuration of the Warehouse Proxy Agent for tuning performance.	“Tuning the performance of the Warehouse Proxy” on page 666
Start the Warehouse Proxy Agent.	“Starting the Warehouse Proxy” on page 632

Creating a TNS Service Name

Create a TNS (Transparent Network Substrate) Service Name (also called a Net Service Name) on a computer where an Oracle client is installed if the Tivoli Data Warehouse exists on a *remote* Oracle server.

Before you begin

This procedure uses the default value for the warehouse name (WAREHOU). Substitute a different value if you do not want to use the default name.

About this task

The TNS Service name is needed to create an ODBC connection between the client and the server. Use this procedure to create a TNS Service name:

- On a Windows computer where a Warehouse Proxy Agent installed.
- On a Windows computer where the Tivoli Enterprise Portal Server is installed.

Do *not* perform this procedure on the computer where the data warehouse (Oracle server) is installed or on a computer where there is no Oracle client (for example, on a computer where a Type 4 Oracle JDBC driver is used to communicate with the remote data warehouse).

Complete the following steps to create the TNS Service Name. Click **Next** after each step.

Procedure

1. Enter dbca at the Oracle command-line to start the Oracle Net Configuration Assistant tool.
2. On the Welcome window, select **Local Net Service Name configuration**.
3. Select **Add**.
4. Enter WAREHOU in the **Service Name** field. (This is the remote name for the Tivoli Data Warehouse.)
5. Select **TCP** as the network protocol to communicate with the Tivoli Data Warehouse database.
6. Specify the fully qualified host name and port number of the computer where the warehouse database is installed.
7. Perform the connection test to verify the connection to the warehouse database.
8. Optionally change the default name in the **Net Service Name** field. This is the TNS Service Name. The default name matches the name that you entered in Step 4. You can change this to a different name. The TNS Service Name can be considered a local alias for the remote Tivoli Data Warehouse name.
9. When prompted to configure another net service name, click **No** to return to the Welcome window.
10. Click **Finish**.

Configuring an ODBC data source for an Oracle data warehouse

An Oracle client on Windows requires an ODBC connection to the data warehouse. For the Warehouse Proxy Agent, you must configure the ODBC connection manually.

Before you begin

This procedure uses default values for the data source name (ITM Warehouse) and warehouse user ID (ITMUser). (Default values are used in configuration procedures for warehousing components.) Substitute different values if you do not want to use the default values.

About this task

Complete the following procedure to set up an ODBC connection for a Warehouse Proxy Agent on Windows to a local or remote Tivoli Data Warehouse.

Procedure

1. On the computer where the Warehouse Proxy Agent is installed, open the Control Panel.
2. Click **Administrative Tools** → **Data Sources (ODBC)**
3. Click **Add** in the **System DSN** tab in the ODBC Data Source Administrator window.
4. Select the Oracle ODBC driver:
 - For Oracle 9, the ODBC driver name is **Oracle in Ora9ias_home**.
 - For Oracle 10, the ODBC driver name is **Oracle in OraDb10g_home1**.
5. Enter ITM Warehouse in the **Data Source Name** field.
6. Enter the TNS Service Name in the **TNS Service Name** field. This is the name that you specified in Step 8 on page 626 (for example, WAREHOUS).
7. Enter ITMUser in the **User ID** field.
8. Click **Test Connection**.
9. In the Oracle ODBC Driver Connect window, enter the TNS Service Name, and the user ID and password of the warehouse user.
10. Click **OK** on the **Connection successful** message.

Configuring a Warehouse Proxy Agent on Windows (ODBC connection)

Use this procedure to configure a Warehouse Proxy Agent on Windows to connect to an Oracle Tivoli Data Warehouse.

Procedure

1. Log on to the Windows system where the Warehouse Proxy Agent is installed and begin the configuration:
 - a. Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Monitoring Services**.
The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Warehouse Proxy** and click **Configure Using Defaults**.
Click **Reconfigure** if the Warehouse Proxy is installed on the same computer as the portal server.
 - c. Click **OK** on the message regarding connection to a hub monitoring server.
2. The next two windows (entitled Warehouse Proxy: Agent Advanced Configuration) contain the settings for the connection between the Warehouse Proxy Agent and the hub monitoring server. Click **OK** on each window to accept the settings.
3. Click **Yes** on the message asking if you want to configure the ODBC data source.
4. Select **Oracle** from the list of databases and click **OK**.

The following configuration window is displayed.

Figure 133. Configure Oracle Data Source for Warehouse Proxy window

5. Click **OK** to accept all default information on this window, or change one or more default values and then click **OK**. The fields on this window are described in Table 114.

Note: The values for the data source name, database name, and database user ID and password must match the values that you used when configuring an ODBC connection for the Warehouse Proxy Agent. (See “Configuring an ODBC data source for an Oracle data warehouse” on page 627.)

Table 114. Configuration information for the Tivoli Data Warehouse database on Oracle

Field	Default value	Description
ODBC DSN	ITM Warehouse	The name of the data source.
Username	ITMUser	The name of the user that the Warehouse Proxy Agent will use to access the Tivoli Data Warehouse database.
Password	itmpswd1	The password that the Warehouse Proxy Agent will use to access the Tivoli Data Warehouse database. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Confirm Password	itmpswd1	Confirm the password by entering it again.
Test Connection		Test the connection to the Tivoli Data Warehouse Database based on the completed fields above: ODBC DSN, Username, and Password. Note: Test Connection is not available if configuring remotely from the Tivoli Enterprise Portal.

Table 114. Configuration information for the Tivoli Data Warehouse database on Oracle (continued)

Field	Default value	Description
Warehouse TEMS List		Environment variable containing a space delimited list of TEMS names, which are given during the configuration of HTEMS or RTEMS. A TEMS name in this field indicates that all the agents connected to this TEMS will have their historical data sent to this Warehouse Proxy Agent. This variable is used when the ITM environment contains multiple Warehouse Proxy Agents and the workload has to be balanced using specific Warehouse Proxy Agents.
Use Batch		Batch inserts can greatly increase the data insertion rate of the Warehouse Proxy Agent. This is especially true if the proxy and the warehouse are located on different hosts. Batch inserts are supported for ODBC warehouse connections. Using batch inserts is recommended in all configurations, but they will place increased load on the data warehouse.
Database Compression		If the database compression mode is supported by the database, the Warehouse Proxy Agent will create all tables and indexes in the Tivoli Data Warehouse with compression enabled. This option reduces the storage costs of the Tivoli Data Warehouse.
Warehouse Compression for Z/OS Sources		Select this option for the Warehouse Proxy server to allow clients installed on Z/OS machines to send compressed data.
Warehouse Compression for Distributed Sources		Select this option for the Warehouse Proxy server to allow clients installed on distributed machines (Linux/UNIX, Windows) to send compressed data.

6. Click OK.

Configuring a Warehouse Proxy Agent on Linux or UNIX (JDBC connection)

Use this procedure to configure a Warehouse Proxy Agent on Linux or UNIX to connect to an Oracle data warehouse.

Procedure

1. Log on to the computer where the Warehouse Proxy Agent is installed and begin the configuration.
 - a. Change to the *install_dir*/bin directory and run the following command:


```
./itmcmd manage [-h install_dir]
```

where *install_dir* is the installation directory for IBM Tivoli Monitoring. The default installation directory is /opt/IBM/ITM.

The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Warehouse Proxy** and click **Configure**.
The Configure Warehouse Proxy window is displayed.




Figure 134. Configure Warehouse Proxy window (Database Type pane)

2. In the **Database** drop-down list, select **Oracle**.
3. Select the **Agent Parameters** check box.




Figure 135. Configure Warehouse Proxy window (Agent Parameters pane)

4. Add the name and directory location of the JDBC driver JAR file to the **JDBC JARS** list box:
 - a. Click **Add** to display the file browser window. Navigate to the location of the JDBC driver JAR file on this computer and select the file. The Oracle JDBC driver JAR file name and default location after downloading from the Web are as follows:

`oracleinstall\dir\jdbc\lib\ojdbc14.jar`

- b. Click **OK** to close the browser window and add the driver file to the list.

If you need to delete an entry from the list, select the entry and click **Delete**.

5. Change the default value displayed in the **JDBC URL** field if it is not correct. The default Tivoli Data Warehouse URL for Oracle is as follows:

`jdbc:oracle:thin:@localhost:1521:WAREHOUS`

- If the Tivoli Data Warehouse is installed on a remote computer, specify the host name of the remote computer instead of `localhost`.
- Change the port number if it is different.
- If the name of the Tivoli Data Warehouse database is not `WAREHOUS`, replace `WAREHOUS` with the actual name. (See “Creating the warehouse database on Oracle” on page 623.)

6. Verify the JDBC driver name, which is displayed in the **JDBC Driver** field. (Note that the **JDBC Driver** field displays the *driver name*, in contrast to the *driver JAR file* that is listed in the **JDBC JARS** field.)

The Oracle JDBC Type 4 driver name is as follows:

`oracle.jdbc.driver.OracleDriver`

7. If necessary, change the entries in the **Username** and **Password** fields to match the user name and password that were created for the Tivoli Data Warehouse. (See “Creating the warehouse database on Oracle” on page 623.) The default user name is `itmuser` and the default password is `itmpswd1`.

8. Check the **Use Batch** check box if you want the Warehouse Proxy Agent to submit multiple execute statements to the Tivoli Data Warehouse database for processing as a batch.
In some situations, such as crossing a network, sending multiple statements as a unit is more efficient than sending each statement separately. Batch processing is one of the features provided with the JDBC 2.0 API.
9. Click **Test connection** to ensure you can communicate with the Tivoli Data Warehouse database.
10. Click **Save** to save your settings and close the window.

Starting the Warehouse Proxy

This section contains instructions for starting the Warehouse Proxy Agent.

- To start the Warehouse Proxy Agent from the Manage Tivoli Enterprise Monitoring Services window, right-click **Warehouse Proxy** and select **Start**.
- (*Linux or AIX only*) To start the Warehouse Proxy Agent from the command-line, run the following command from the bin directory of the IBM Tivoli Monitoring installation directory. The default installation directory is /opt/IBM/ITM.
`./itmcmd agent start hd`

where hd is the product code for the Warehouse Proxy Agent.

Step 3: Configure communications between the Tivoli Enterprise Portal Server and the data warehouse

Complete the tasks described in the following table, in the order listed, to configure communications between the portal server and the data warehouse.

Table 115. Tasks for configuring communications between the portal server and an Oracle data warehouse

Task	Procedure
(<i>Portal server on Windows only</i>) Install an Oracle database client on the portal server.	See the Oracle documentation for instructions on how to install an Oracle client.
(<i>Portal server on Windows only</i>) Create a TNS (Transparent Network Substrate) Service Name on the portal server.	"Creating a TNS Service Name" on page 626
(<i>Portal server on Linux or AIX only</i>) Install an Oracle JDBC Type 4 driver on the portal server.	Obtain the Oracle JDBC Type 4 Driver from the following Web site: http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html The Oracle JDBC driver JAR file name and location after installation is as follows: <i>oracleinstalldir/jdbc/lib/ojdbc14.jar</i> The ojdbc14.jar file supports JRE 1.5 or higher, the required Java Runtime Environment for IBM Tivoli Monitoring.
Configure the portal server to connect to the data warehouse. The configuration procedure on Windows automatically configures an ODBC connection to the data warehouse.	For a portal server on Windows, see "Configuring a Windows portal server (ODBC connection)" on page 633. For a portal server on Linux or AIX, see "Configuring a Linux or AIX portal server (JDBC connection)" on page 634.

Table 115. Tasks for configuring communications between the portal server and an Oracle data warehouse (continued)

Task	Procedure
Restart the portal server.	"Starting the portal server" on page 636
Test the connection between the portal server and the Tivoli Data Warehouse by creating a customized query in the Tivoli Enterprise Portal.	"Testing the connection between the portal server and the Tivoli Data Warehouse" on page 663

Configuring a Windows portal server (ODBC connection)

Use this procedure to configure a portal server on Windows to connect to an Oracle Tivoli Data Warehouse.

Procedure

- Log on to the Windows system where the portal server is installed and begin the configuration:
 - Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Monitoring Services**.
The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - Right-click **Tivoli Enterprise Portal Server** and click **Reconfigure**.
- The next two windows (entitled TEP Server Configuration) contain the settings for the connection between the portal server and the hub monitoring server. These settings were specified when the portal server was installed. Click **OK** on each window to accept the settings.
- Click **Yes** on the message asking if you want to reconfigure the warehouse information for the Tivoli Enterprise Portal Server.
- Select **Oracle** from the list of databases and click **OK**.
The following configuration window is displayed.

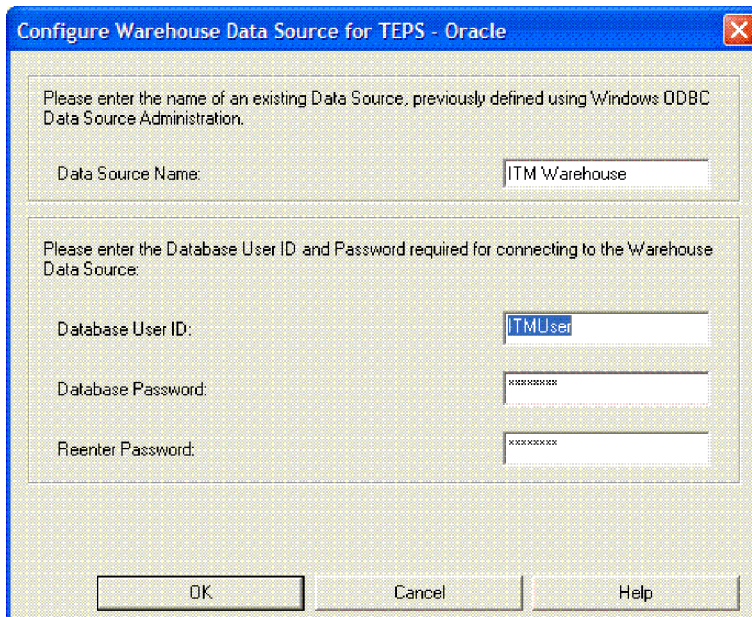


Figure 136. Configure Oracle Data Source for Warehouse window

- Click **OK** to accept all default information on this window, or change one or more default values and then click **OK**. The fields on this window are described in the following table:

Table 116. Configuration information for the Tivoli Data Warehouse database on Oracle

Field	Default value	Description
Data Source Name	ITM Warehouse	The name of the data source.
Database User ID	ITMUser	The login name of the database user that the portal server will use to access the Tivoli Data Warehouse database.
Database Password	itmpswd1	The password for the database login user. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Reenter Password	itmpswd1	Confirm the password by entering it again.

6. Click **OK**.

Configuring a Linux or AIX portal server (JDBC connection)

Use this procedure to configure a portal server on Linux or AIX to connect to an Oracle Tivoli Data Warehouse on any operating system.

Procedure

1. Log on to the computer where the Tivoli Enterprise Portal Server is installed and begin the configuration.
 - a. Change to the *install_dir/bin* directory and run the following command:


```
./itmcmd manage [-h install_dir]
```

where *install_dir* is the installation directory for IBM Tivoli Monitoring. The default installation directory is */opt/IBM/ITM*.

The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Tivoli Enterprise Portal Server** and click **Configure**.

The Configure Tivoli Enterprise Portal Server window is displayed.
2. On the **TEMS Connection** tab, review the settings for the connection between the portal server and the hub monitoring server. These settings were specified when the portal server was installed.
3. Click the **Agent Parameters** tab.
4. Select the **Oracle** radio button.

The fields for configuring the connection to an Oracle data warehouse are displayed at the bottom of the window.

Figure 137. Configuring the connection to an Oracle data warehouse

5. Fill in the fields in Figure 137 with the configuration values described in Table 117.

Table 117. Configuration information for a Tivoli Data Warehouse database on Oracle

Field	Default value	Description
Warehouse database name	WAREHOUS	The name of the Tivoli Data Warehouse database.
Warehouse DB user ID	ITMUser	The database login user that the portal server uses to access the Tivoli Data Warehouse database. This user is referred to as the <i>warehouse user</i> .
Warehouse user password	itmpswd1	The password for the warehouse user.
Re-type Warehouse user password	itmpswd1	The password for the warehouse user.
JDBC driver class path	<i>oracleinstalldir/jdbc/lib/ojdbc14.jar</i>	The full path name of the Oracle JDBC Type 4 driver JAR file on this computer.
JDBC driver name	oracle.jdbc.driver.OracleDriver	The Oracle JDBC Type 4 driver name.
JDBC driver URL	jdbc:oracle:thin:@localhost:1521:WAREHOUS	<p>The Oracle-defined URL that identifies the Oracle instance used for the remote Tivoli Data Warehouse.</p> <p>Replace <i>localhost</i> with the host name of the remote computer where the Tivoli Data Warehouse is installed.</p> <p>Change the default port number (1521) and Tivoli Data Warehouse name (WAREHOUS) if they are different.</p>
User-defined attributes	(no default)	Enter any user-defined attributes that are used to customize the behavior of the driver connection. Use semi-colons (;) to delimit the attributes.

6. Click **Save** to save your settings and close the window.

Starting the portal server

This section contains instructions for starting the portal server.

- To start the portal server from the Manage Tivoli Enterprise Monitoring Services window, right-click **Tivoli Enterprise Portal Server** and select **Start**.
- (*Linux or AIX only*) To start the portal server from the command-line, run the following command from the bin directory of the IBM Tivoli Monitoring installation directory. The default installation directory is /opt/IBM/ITM.

```
./itmcmd agent start cq
```

where cq is the product code for the portal server.

Step 4: Install and configure communications for the Summarization and Pruning Agent

Complete the tasks described in this table, in the order listed, to install and configure the Summarization and Pruning Agent.

Table 118. Tasks for installing and configuring communications for the Summarization and Pruning Agent

Task	Procedure
<p>Install the Summarization and Pruning Agent if you have not already installed it. For best performance, install the Summarization and Pruning Agent on the same computer as the data warehouse.</p> <p>The installation procedure for Windows includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring server. On Linux or AIX, this step is performed in a separate configuration procedure (<i>Configuring the monitoring agent</i>). See the information at right. Be sure to perform all referenced installation and configuration procedures.</p> <p>Note: The Summarization and Pruning Agent is not automatically started after installation. Do not complete any step or procedure for starting the agent at this point.</p>	<p>To install a Summarization and Pruning Agent on Windows, complete the procedure “Windows: Installing a monitoring agent” on page 262.</p> <p>To install a Summarization and Pruning Agent on Linux or UNIX, complete the procedure “Linux or UNIX: Installing a monitoring agent” on page 267, including the following subsections:</p> <ul style="list-style-type: none">• <i>Installing the monitoring agent</i>• <i>Configuring the monitoring agent</i>• <i>Changing the file permissions for agents</i> (if you used a non-root user to install the Warehouse Proxy) <p><i>Do not complete the procedure for starting the agent.</i></p>
<p>Install an Oracle Type 4 JDBC driver on the computer where the Summarization and Pruning Agent is installed.</p>	<p>Obtain the Oracle JDBC Driver from the following Web site:</p> <p>http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html</p> <p>The Oracle JDBC driver JAR file name and location after installation is as follows:<i>oracleinstalldir/jdbc/lib/ojdbc14.jar</i></p> <p>The ojdbc14.jar file supports JRE 1.5 or higher, the required Java Runtime Environment for IBM Tivoli Monitoring.</p>

Table 118. Tasks for installing and configuring communications for the Summarization and Pruning Agent (continued)

Task	Procedure
<p>Configure the Summarization and Pruning Agent.</p> <p>When you configure the Summarization and Pruning Agent, you configure the connection to the Tivoli Data Warehouse and you specify settings that control the operation of the Summarization and Pruning Agent.</p> <p>Perform this procedure whether or not the Summarization and Pruning Agent and the warehouse database are installed on the same computer.</p>	<p>“Configuring the Summarization and Pruning Agent (JDBC connection)” on page 643</p>
<p>Configure the Summarization and Pruning Agent to connect to the Tivoli Enterprise Portal Server. Perform this procedure whether or not the Summarization and Pruning Agent and the warehouse database are installed on the same computer.</p>	<p>See Step 9 of “Configuring the Summarization and Pruning Agent (JDBC connection)” on page 643.</p>
<p>Configure history collection.</p> <p>When you configure history collection, you specify settings for how often to collect, aggregate, and prune data for individual monitoring agents and attribute groups. Configure history collection from the Tivoli Enterprise Portal.</p>	<p>See the <i>IBM Tivoli Monitoring: Administrator's Guide</i> for instructions on how to configure history collection.</p>
<p>Start the Summarization and Pruning Agent.</p>	<p>“Starting the Summarization and Pruning Agent” on page 656</p>

Step 5: Install and configure communications for Tivoli Performance Analyzer

You can install Tivoli Performance Analyzer to a server which also has other Tivoli Monitoring components installed, or you can install it to a separate machine. The installation procedure is similar to that for monitoring agents.

Complete the tasks described in the following table, in the order listed, to install and configure Tivoli Performance Analyzer.

Table 119. Tasks for installing and configuring communications for the Tivoli Performance Analyzer

Task	Procedure
<p>Install Tivoli Performance Analyzer.</p> <p>The installation procedure for Windows includes steps for configuring the connection between the Tivoli Performance Analyzer and the hub Tivoli Enterprise Monitoring server. On Linux or AIX, this step is performed in a separate configuration procedure (Configuring the monitoring agent). See the information opposite. Be sure to perform all of the referenced installation and configuration procedures.</p> <p>Note: for sites setting up autonomous operation, the installation procedure includes steps for configuring the connection between the agent and the hub Tivoli Enterprise Monitoring Server. On Windows operating systems, if you want to run Tivoli Performance Analyzer without a connection to the hub, accept the defaults for the connection information, but specify a non valid name for the monitoring server. On UNIX and Linux operating systems, select No TEMS on the TEMS Connection in the configuration window.</p>	<p>To install Tivoli Performance Analyzer on Windows, complete the procedure “Windows: Installing a monitoring agent” on page 262. To install Tivoli Performance Analyzer agent on Linux or AIX, complete the procedure “Linux or UNIX: Installing a monitoring agent” on page 267, including the following subsections:</p> <ul style="list-style-type: none"> • Installing the monitoring agent • Configuring the monitoring agent • Changing the file permissions for agents <p>Do not complete the procedure for starting the agent.</p>

Table 119. Tasks for installing and configuring communications for the Tivoli Performance Analyzer (continued)

Task	Procedure
<p>Ensure that the latest Oracle patches are installed. Set the following system variable on the computer where Tivoli Performance Analyzer is installed. Restart the computer after setting the variable. The format of the NLS_LANG environment variable is: NLS_LANG=language_territory.charset.</p> <p>Set the language and territory to appropriate variables. For the United States this is NLS_LANG=AMERICAN_AMERICA.AL32UTF8.</p>	<p>See the Oracle documentation for instructions on how to install an Oracle client. Obtain Oracle ODBC drivers from the following Web site:http://www.oracle.com/technology/software/tech/windows/odbc/htdocs/utilsoft.html.</p>
<p>If the Tivoli Data Warehouse is located on a remote computer, create a TNS (Transparent Network Substrate) Service Name on the local computer where the Tivoli Performance Analyzer is located.</p>	<p>“Creating a TNS Service Name”</p>
<p>On the computer where the Tivoli Performance Analyzer is installed, configure an ODBC data source for the data warehouse, using the TNS Service Name that you created in the preceding step.</p>	<p>“Configuring an ODBC data source for an Oracle data warehouse” on page 639</p>
<p>Install an Oracle Type 4 JDBC driver on the computer where the Tivoli Performance Analyzer is installed.</p>	<p>Obtain the Oracle JDBC Driver from the following Web site:</p> <p>http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html</p> <p>The Oracle JDBC driver JAR file name and location after installation is as follows:</p> <p><i>oracleinstall\dir\jdbc\lib\ojdbc14.jar</i></p> <p>The ojdbc14.jar file supports JRE 1.5 or higher, the required Java Runtime Environment for IBM Tivoli Monitoring.</p>
<p>Configure Tivoli Performance Analyzer to connect to the data warehouse.</p>	<p>For Tivoli Performance Analyzer on Windows, see “Configuring a Tivoli Performance Analyzer on Windows (ODBC connection)” on page 640. For a Tivoli Performance Analyzer on Linux or AIX, see “Configuring Tivoli Performance Analyzer on Linux or UNIX (JDBC connection)” on page 641.</p>
<p>Start Tivoli Performance Analyzer.</p>	<p>“Starting Tivoli Performance Analyzer” on page 642</p>

Creating a TNS Service Name

Create a TNS (Transparent Network Substrate) Service Name (also called a Net Service Name) on a computer where an Oracle client is installed if the Tivoli Data Warehouse exists on a *remote* Oracle server. The TNS Service name is needed to create an ODBC connection between the client and the server.

Before you begin

This procedure uses the default value for the warehouse name (WAREH0US). Substitute a different value if you do not want to use the default name. Do not perform this procedure on the computer where the data warehouse (Oracle server) is installed or on a computer where there is no Oracle client (for example, on a computer where a Type 4 Oracle JDBC driver is used to communicate with the remote data warehouse).

About this task

Use this procedure to create a TNS Service name:

- On a Windows computer where Tivoli Performance Analyzer is installed.
- On a Windows computer where the Tivoli Enterprise Portal Server is installed.

Complete the following steps to create the TNS Service Name. Click **Next** after each step.

Procedure

1. Enter **dbca** at the Oracle command-line to start the Oracle Net Configuration Assistant tool.
2. On the Welcome window, select **Local Net Service Name configuration**.
3. Select **Add**.
4. Enter **WAREHOU** in the **Service Name** field. (This is the remote name for the Tivoli Data Warehouse.)
5. Select **TCP** as the network protocol to communicate with the Tivoli Data Warehouse database.
6. Specify the fully qualified host name and port number of the computer where the warehouse database is installed.
7. Perform the connection test to verify the connection to the warehouse database.
8. Optionally change the default name in the **Net Service Name** field.
This is the TNS Service Name. The default name matches the name that you entered in Step 4 on page 626. You can change this to a different name. The TNS Service Name can be considered a local alias for the remote Tivoli Data Warehouse name.
9. When prompted to configure another net service name, click **No** to return to the Welcome window.
10. Click **Finish**.

Configuring an ODBC data source for an Oracle data warehouse

An Oracle client on Windows requires an ODBC connection to the data warehouse. For the Tivoli Performance Analyzer, you must configure the ODBC connection manually.

Before you begin

This procedure uses default values for the data source name (ITM Warehouse) and warehouse user ID (ITMUser). (Default values are used in configuration procedures for warehousing components.) Substitute different values if you do not want to use the default values.

About this task

Complete the following procedure to set up an ODBC connection for a Tivoli Performance Analyzer on Windows to a local or remote Tivoli Data Warehouse.

Procedure

1. On the computer where the Tivoli Performance Analyzer is installed, open the Control Panel.
2. Click **Administrative Tools → Data Sources (ODBC)**
3. Click **Add** in the **System DSN** tab in the ODBC Data Source Administrator window.
4. Select the Oracle ODBC driver:
 - For Oracle 9, the ODBC driver name is **Oracle in Ora9ias_home**.
 - For Oracle 10, the ODBC driver name is **Oracle in OraDb10g_home1**.
5. Enter **ITM Warehouse** in the **Data Source Name** field.
6. Enter the TNS Service Name in the **TNS Service Name** field. This is the name that you specified in Step 8 on page 626 (for example, **WAREHOU**).
7. Enter **ITMUser** in the **User ID** field.

8. Click **Test Connection**.
9. In the Oracle ODBC Driver Connect window, enter the TNS Service Name, and the user ID and password of the warehouse user.
10. Click **OK** on the **Connection successful** message.

Configuring a Tivoli Performance Analyzer on Windows (ODBC connection)

Use this procedure to configure Tivoli Performance Analyzer on Windows to connect to an Oracle Tivoli Data Warehouse.

Procedure

1. Log on to the Windows system where the Tivoli Performance Analyzer is installed and begin the configuration:
 - a. Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Monitoring Services**. The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Tivoli Performance Analyzer** and click **Reconfigure**.
 - c. Click **OK** on the message regarding connection to a hub monitoring server.
2. The next two windows (entitled Performance Analyzer: Agent Advanced Configuration) contain the settings for the connection between the Tivoli Performance Analyzer and the hub monitoring server. These settings were specified when Tivoli Performance Analyzer was installed. Click **OK** on each window to accept the settings.
3. Click **Yes** on the message asking if you want to configure the ODBC data source.
4. Select **ODBC** from the list of selectable agent database connection types.
5. Set your Database Type to Oracle.
6. Specify the Data Source Name - **Agent ODBC DSN (ITM Warehouse by default)**.

Note:

- Tivoli Performance Analyzer does not create this DSN - it must already exist. If you are installing Performance Analyzer on the same machine where TEP Server is installed, you can use the existing data source created by Tivoli Monitoring. Otherwise, you must create a new System DSN manually, prior to re-configuring Performance Analyzer.
 - On 64-bit versions of Windows, data sources created by the default **ODBC Data Source Administrator** applet available from the **Control Panel** are not available for 32-bit applications. Therefore you must use the 32-bit version of the **ODBC Data Source Administrator** applet from <WINDOWS>\SysWOW64\odbcad32.exe.
7. Type the **Username** and **Password**. The entries in these fields are used to connect to the Tivoli Data Warehouse and are the same credentials as those used by the Tivoli Enterprise Portal Server, the Warehouse Proxy Agent and the Summarization and Pruning Agent to communicate with Tivoli Data Warehouse.
 8. Click **Next** to proceed to the Advanced Configuration window.
 9. You can enable Advanced Configuration to specify TDW Schema and The TDW database schema. If you do not select **Enable advanced configuration** these options are greyed out.
 10. You can also choose whether you want the agent to Initialize PA tables and OS domain tasks.

Note: Setting Initialize PA tables to YES will remove and recreate all previously created tables deleting all user tasks and reverting each OS task to its default.

11. Use the **Bypass connection tests** option to finish the configuration without running connection tests.
12. Click **OK** to finish the configuration process.

Note: The values for the data source name, and database user ID and password must match the values that you used when configuring an ODBC connection for Tivoli Performance Analyzer.

Table 120. Configuration information for the Tivoli Data Warehouse database on Oracle

Field	Default value	Description
ODBC DSN	ITM Warehouse	The name of the data source.
Username	ITMUser	The name of the Windows OS user that the Tivoli Performance Analyzer will use to access the Tivoli Data Warehouse database.
Password	itmpswd1	The password for the Windows OS user. If your environment requires complex passwords (passwords that require both alpha and numeric characters), specify a password that complies with these requirements.
Test Connection		Test the connection to the Tivoli Data Warehouse Database based on the completed fields above: ODBC DSN, Username, and Password.

Configuring Tivoli Performance Analyzer on Linux or UNIX (JDBC connection)

Use this procedure to configure Tivoli Performance Analyzer on Linux or UNIX to connect to an Oracle data warehouse.

Procedure

1. To begin the configuration, log on to the computer where Tivoli Performance Analyzer is installed.
 - a. Change to the `install_dir/bin` directory and run the following command:

```
./itmcmd manage [-h install_dir]
```

where `install_dir` is the installation directory for IBM Tivoli Monitoring. The default installation directory is `/opt/IBM/ITM`. The Manage Tivoli Enterprise Monitoring Services window is displayed.
 - b. Right-click **Performance Analyzer** and click **Configure**. The Configure Tivoli Performance Analyzer window is displayed.
2. Set the Database Type to Oracle.
3. Type the username and the password. The entries in these fields are used to connect to the Tivoli Data Warehouse.
4. Review all the defaults in the Agent Configuration window and change as required.
 - a. If the Tivoli Data Warehouse is installed on a remote computer, specify the host name of the remote computer instead of `localhost`.
 - b. Change the port number if necessary (the default port number for Oracle is 1521).
 - c. If the name of the Tivoli Data Warehouse database is not `WAREHOUS`, replace `WAREHOUS` with the actual name. (See “Creating the warehouse database on DB2 for Linux, UNIX and Windows” on page 544.)
5. Specify the JDBC Driver. The Oracle JDBC Type 4 driver name is `oracle.jdbc.driver.OracleDriver`.
6. Specify the JDBC Driver Path, which should be provided as a list of JAR files with the full path separated by `“:”`.

Note: The Oracle JDBC driver JAR file name and default location after downloading from the Web are as follows:

`doracleinstall_dir/jdbc/lib/ojdbc14.jar`

Fast path: You can use the Browse button to specify the path. In such a case a file list is added at the end of the JDBC Driver Path text field, separated from the existing content by a path separator.

7. You can use the **Test connection** button to check whether the connection can be initiated.
8. Click **Next** to proceed to the Advanced Configuration window.
 - a. You can enable Advanced Configuration to specify TDW Schema and Configuration schema. If you do not select Enable advanced configuration, all of these options are greyed out.
 - b. You can also choose whether you want the agent to initialize PA tables.

Note: Setting Initialize PA tables to YES will remove and recreate all previously created tables deleting all user tasks and reverting each OS task to its default.

- c. Use the **Bypass connection tests** option to finish the configuration without running connection tests.
9. Click **Save** to save your settings and close the window.

Starting Tivoli Performance Analyzer

This section contains instructions for Starting the Tivoli Performance Analyzer.

To start Tivoli Performance Analyzer from the Manage Tivoli Enterprise Monitoring Services window, right-click Tivoli Performance Analyzer and select Start. To start the Tivoli Performance Analyzer agent from the command-line, run the following command from the bin directory of the IBM Tivoli Monitoring installation directory. The default installation directory is /opt/IBM/ITM.

```
./itmcmd agent start pa
```

where pa is the product code for Tivoli Performance Analyzer agent.

Chapter 27. Tivoli Data Warehouse solutions: common procedures

This chapter contains information and procedures for Tivoli Data Warehouse solutions that use any of the supported database platforms: IBM DB2 for Linux, UNIX, and Windows, Microsoft SQL Server, and Oracle.

Read the chapter pertaining to the database platform you are using for the Tivoli Data Warehouse before performing any of these procedures.

- “Configuring the Summarization and Pruning Agent (JDBC connection)”
- “Starting the Summarization and Pruning Agent” on page 656
- “Installing and configuring multiple Warehouse Proxy Agents” on page 656
- “Running the warehouse agents autonomously” on page 658
- “Testing the connection between the portal server and the Tivoli Data Warehouse” on page 663
- “Tuning the performance of the Warehouse Proxy” on page 666
- “WAREHOUSELOG and WAREHOUSEAGGREGLOG tables” on page 689
- “Configuring the Warehouse Proxy Agent on Linux or UNIX: command-line procedure” on page 653
- “Configuring the Summarization and Pruning Agent on Linux or UNIX: command-line procedure” on page 654

Configuring the Summarization and Pruning Agent (JDBC connection)

Use this procedure to configure the Summarization and Pruning Agent to connect to a Tivoli Data Warehouse database created on any of the supported database platforms and operating systems.

Before you begin

The JDBC driver JAR files for your database platform must be located on the computer where you installed the Summarization and Pruning Agent. Use a Type 4 JDBC driver. Do not use the Type 2 driver.

- If you are using DB2 for Linux, UNIX, and Windows for your Tivoli Data Warehouse database, the JDBC driver files are included with the database platform installation. If the warehouse database and the Summarization and Pruning Agent are installed on the same computer, the JDBC driver files are already present. It is not necessary to move them from their current location. If your warehouse database is located on a remote computer, copy the driver files to the local computer (the computer where you installed the Summarization and Pruning Agent). Copy the files to any directory to which the user the Summarization and Pruning Agent is running as has access to.
- If you are using Oracle or Microsoft SQL Server for your Tivoli Data Warehouse database, download the driver file from the company Web site to the computer where you installed the Summarization and Pruning Agent. Download the files to any directory to which the user the Summarization and Pruning Agent is running as has access to.

Table 121 on page 644 shows where to obtain the driver files for each database platform.

Table 121. Where to obtain the JDBC driver files for the Summarization and Pruning Agent

Database platform	JDBC driver files
IBM DB2 for Linux, Unix and Windows	<p>Use the DB2 for Linux, UNIX, and Windows JDBC Universal Driver (Type 4 driver). The DB2 for Linux, UNIX, and Windows driver files are located with your Tivoli Data Warehouse server installation. The Type 4 driver file names and locations are as follows:</p> <p><i>db2installdir</i>/java/db2jcc.jar <i>db2installdir</i>/java/db2jcc_license_cu.jar</p> <p>where <i>db2installdir</i> is the directory where DB2 for Linux, UNIX, and Windows was installed. The default DB2 for Linux, UNIX, and Windows Version 9 installation directory is as follows:</p> <ul style="list-style-type: none"> • On Windows: C:\Program Files\IBM\SQLLIB • On AIX: /usr/opt/db2_09_01 • On Linux and Solaris: /opt/IBM/db2/V9.1
Microsoft SQL Server	<p>Use the latest Microsoft SQL Server JDBC driver, which supports SQL Server 2008, 2005, and 2000, to connect to a Tivoli Data Warehouse on either SQL Server 2000 or SQL Server 2005. (The SQL Server 2005 JDBC Driver works with a Tivoli Data Warehouse on SQL Server 2000.) Go to the Microsoft Web page at:</p> <p>http://www.microsoft.com and search for JDBC driver.</p> <p>Download and install the driver to the computer where you installed the Summarization and Pruning Agent. Follow the instructions on the Microsoft download page for installing the driver. The SQL Server JAR file name and location after installation is as follows:<<i>mssqlinstall</i>dir>/sqljdbc_1.1/enu/sqljdbc4.jar.</p>
Oracle	<p>Obtain the Oracle JDBC Type 4 driver from the following Web site:</p> <p>http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html.</p> <p>The Oracle JDBC driver JAR file name and location after installation is as follows:<i>oracleinstall</i>dir/jdbc/lib/ojdbc14.jar</p> <p>The ojdbc14.jar file supports JRE 1.5 or higher, the required Java Runtime Environment for IBM Tivoli Monitoring.</p>

About this task

Complete the following steps to configure the Summarization and Pruning Agent.

Procedure

- Log on to the computer where the Summarization and Pruning Agent is installed and begin the configuration:
 - Open the Manage Tivoli Enterprise Monitoring Services window:
 - Windows** Click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Monitoring Services**.
 - Linux** **UNIX** Change to the *install_dir*/bin directory and run the following command:
./itmcmd manage [-h *install_dir*]

where *install_dir* is the installation directory for IBM Tivoli Monitoring. The default installation directory is /opt/IBM/ITM.
 - Right-click **Summarization and Pruning Agent**.
 - Windows** Click **Configure Using Defaults**. **Linux** **UNIX** Click **Configure**. If you are reconfiguring, click **Reconfigure**.

2. Review the settings for the connection between the Summarization and Pruning Agent and the hub Tivoli Enterprise Monitoring server. These settings were specified when the Summarization and Pruning Agent was installed.

- **Windows** Perform the following steps:
 - a. On the Warehouse Summarization and Pruning Agent: Agent Advanced Configuration window, verify the communications protocol of the hub monitoring server in the **Protocol** drop down list. Click **OK**.
 - b. On the next window, verify the host name and port number of the hub monitoring server. Click **OK**.

For information about the different protocols available to the hub monitoring server on Windows, and associated default values, see “Windows: Installing the hub monitoring server” on page 196.

- **Linux** **UNIX** Verify the following information on the **TEMS Connection** tab:
 - The hostname of the hub monitoring server in the **TEMS Hostname** field. (If the field is not active, clear the **No TEMS** check box.)
 - The communications protocol that the hub monitoring server uses in the **Protocol** drop down list.
 - If you select IP.UDP, IP.PIPE, or IP.SPIPE, enter the port number of the monitoring server in the **Port Number** field.
 - If you select SNA, enter information in the **Net Name**, **LU Name**, and **LOG Mode** fields.

For information about the different protocols available to the hub monitoring server on Linux or UNIX, and associated default values, see “Configuring the hub monitoring server” on page 202.

3. When you are finished verifying or entering information about the hub monitoring server:
 - **Windows** Click **Yes** on the message asking if you want to configure the Summarization and Pruning Agent.
 - **Linux** **UNIX** Click the **Agent Parameters** tab.

A multi-tabbed configuration window is displayed with the **Sources** tab at the front.

Figure 138 on page 646 shows the configuration window for a Summarization and Pruning Agent on Windows (values displayed are for a DB2 for Linux, UNIX, and Windows warehouse database). The configuration window for a Summarization and Pruning Agent on Linux or UNIX is similar.




Figure 138. Sources pane of Configure Warehouse Summarization and Pruning Agent window

4. Add the names and directory locations of the JDBC driver JAR files to the **JDBC Drivers** list box:
 - a. Click **Add** to display the file browser window. Navigate to the location of the driver files on this computer and select the Type 4 driver files for your database platform. See Table 121 on page 644 for the names and default locations of the driver files to add.
 - b. Click **OK** to close the browser window and add the JDBC driver files to the list.

If you need to delete an entry from the list, select the entry and click **Remove**.
5. The default values for the database platform you selected in the Database Type pane are displayed in the other text fields on the **Sources** pane. Change the default value displayed in the **JDBC URL** field if it is not correct. The following table lists the default Tivoli Data Warehouse URLs for the different database platforms:

Table 122. Tivoli Data Warehouse URLs

Database platform	Warehouse URL
IBM DB2 for Linux, UNIX, and Windows	jdbc:db2://localhost:60000/WAREHOUS
Oracle	jdbc:oracle:thin:@localhost:1521:WAREHOUS
Microsoft SQL Server 2000 or SQL Server 2005	jdbc:sqlserver://localhost:1433;databaseName=WAREHOUS

- If the Tivoli Data Warehouse is installed on a remote computer, specify the host name of the remote computer instead of localhost.
- Change the port number if it is different.
- If the name of the Tivoli Data Warehouse database is not WAREHOUS, replace WAREHOUS with the actual name.

6. Verify the JDBC driver name.

The following table lists the JDBC Type 4 driver names for each database platform:

Table 123. JDBC driver names

Database platform	JDBC driver name
IBM DB2 for Linux, UNIX, and Windows	com.ibm.db2.jcc.DB2Driver
Oracle	oracle.jdbc.driver.OracleDriver
Microsoft SQL Server	com.microsoft.sqlserver.jdbc.SQLServerDriver Note: This is the name of the 2005 SQL Driver. Do not use the SQL Server 2000 JDBC driver, even if the Tivoli Data Warehouse was created in Microsoft SQL 2000. (The name of the 2000 SQL driver was com.microsoft.jdbc.sqlserver.SQLServerDriver. Note the reversal of the string jdbc.sqlserver.)

7. If necessary, change the entries in the **Warehouse user** and **Warehouse password** fields to match the user name and password that were created for the Tivoli Data Warehouse. The default user name is itmuser and the default password is itmpswd1.
8. In the **TEPS Server Host** and **TEPS Server Port** fields, enter the host name of the computer where the Tivoli Enterprise Portal Server is installed and the port number that it uses to communicate with the Summarization and Pruning Agent.

Note: The default Tivoli Enterprise Portal Server interface port of 15001 is also used after the Summarization and Pruning Agent's initial connection to the portal server over port 1920. Any firewalls between the two need to allow communications on either 15001 or whichever port is defined for any new Tivoli Enterprise Portal Server interface used per the instructions in "Defining a Tivoli Enterprise Portal Server interface on Windows" on page 876.

9. Click **Test connection** to ensure you can communicate with the Tivoli Data Warehouse database.

Note: If Microsoft SQL Server is the database platform, an error message might be displayed after clicking **Test connection** even if you have changed the JDBC driver to sqljdbc4.jar. The problem is caused by the caching of the old sqljdbc.jar file. If you save the settings with the new sqljdbc4.jar file and retry **Test connection**, the new jar file is used and the error message will not appear.

10. Select the **Scheduling** check box to specify when you want summarization and pruning to take place. You can schedule it to run on a fixed schedule or on a flexible schedule:

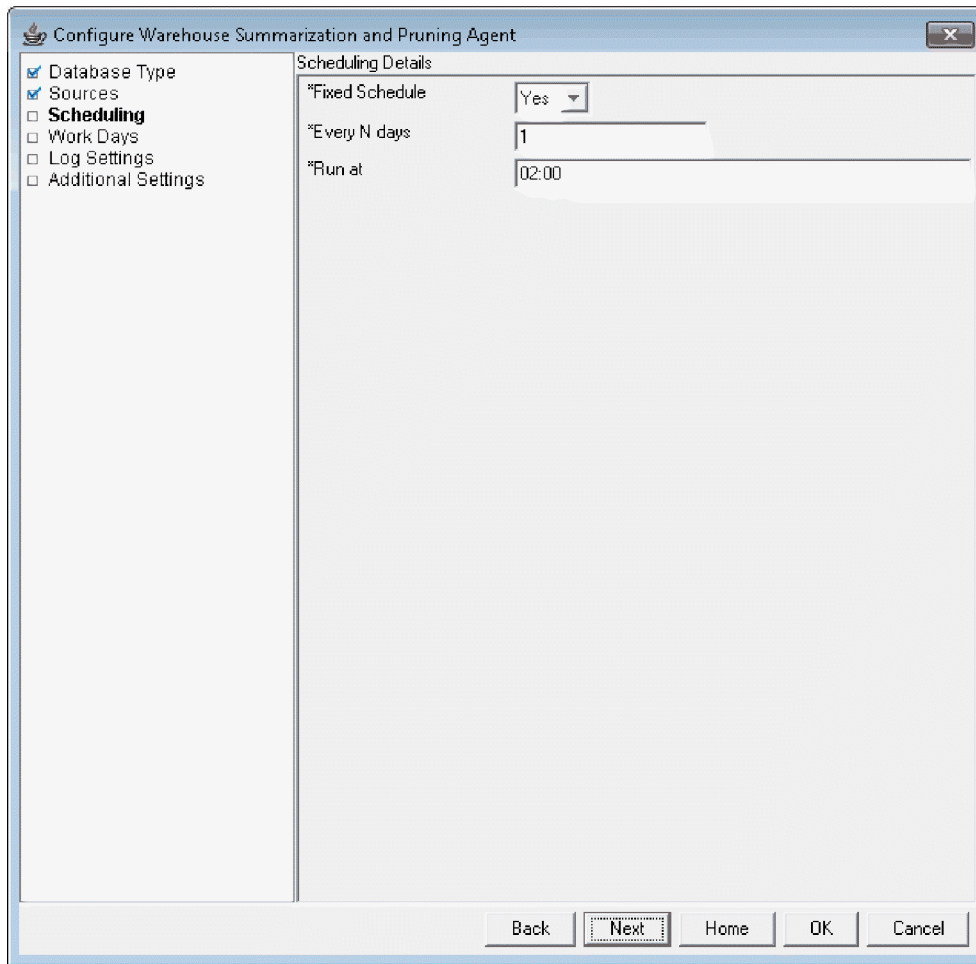


Figure 139. Scheduling pane of Configure Warehouse Summarization and Pruning Agent window

Note: If you select Fixed, the Summarization and Pruning Agent does not immediately perform any summarization or pruning when it *starts*. It performs summarization and pruning when it *runs*. It runs according to the schedule you specify on the **Scheduling** pane. If you select Flexible, the Summarization and Pruning Agent runs once immediately after it is started and then at the interval you specified except during any blackout times.

11. Specify shift and vacation settings in the **Work Days** pane:

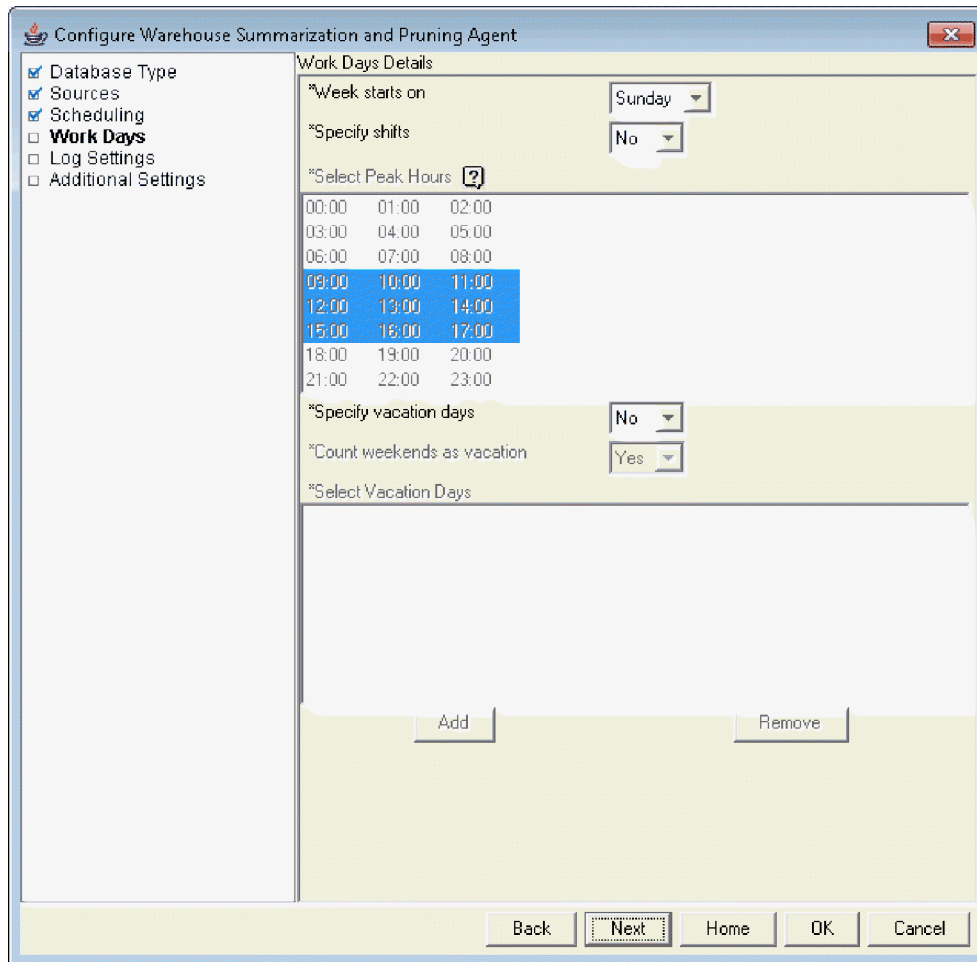


Figure 140. Work Days pane of Configure Warehouse Summarization and Pruning Agent window

When you enable and configure shifts, IBM Tivoli Monitoring produces three separate summarization reports:

- Summarization for peak shift hours
- Summarization for off-peak shift hours
- Summarization for all hours (peak and off-peak)

Similarly, when you enable and configure vacations, IBM Tivoli Monitoring produces three separate summarization reports:

- Summarization for vacation days
- Summarization for nonvacation days
- Summarization for all days (vacation and nonvacation)

Complete the following steps to enable shifts, vacations, or both:

- Select when the beginning of the week starts.
- To configure shifts:
 - a. Select **Yes** in the **Specify shifts** drop-down list.
 - b. Optionally change the default settings for peak and off peak hours by selecting hours in the **Select Peak Hours** box.

Note: Changing the shift information after data has been summarized creates an inconsistency in the data. Data that was previously collected is not summarized again to account for the new shift values.

- To configure vacation settings:
 - a. Select **Yes** in the **Specify vacation days** drop-down list to enable vacation days.
 - b. Select **Yes** in the drop-down list if you want to specify weekends as vacation days.
 - c. Select **Add** to add vacation days.
 - d. Select the vacation days you want to add from the calendar.

On UNIX or Linux, right-click, instead of left-click, to select the month and year.

The days you select are displayed in the list box.

If you want to delete any days you have previously chosen, select them and click **Delete**.

Note:

- 1) Add vacation days in the future. Adding vacation days in the past creates an inconsistency in the data. Data that was previously collected is not summarized again to account for vacation days.
- 2) Enabling shifts or vacation periods can significantly increase the size of the warehouse database. It will also negatively affect the performance of the Summarization and Pruning Agent.

12. Select the **Log Settings** check box to set the intervals for log pruning:




Figure 141. Log Settings pane of Configure Warehouse Summarization and Pruning Agent window

- Select Prune WAREHOUSELOG, select the number of units for which data should be kept, and the unit of time (day, month or year).
- Select Prune WAREHOUSEAGGREGLOG, select the number of units for which data should be kept, and the unit of time (day, month or year).

Note: These table inserts are disabled by default. The self-monitoring workspaces provided by the Summarization and Pruning Agent provide sufficient information to determine if the agent is operating correctly. Tivoli Data Warehouse log tables can grow very large and require regular pruning. This new default configuration decreases the Summarization and Pruning Agent processing time, and decreases database resource utilization and contention.

13. Specify additional summarization and pruning settings in the **Additional Settings** pane:

Figure 142. Additional Settings pane of Configure Warehouse Summarization and Pruning Agent window

- a. Specify the number of additional threads you want to use for handling summarization and pruning processing. The number of threads should be $2 * N$, where N is the number of processors running the Summarization and Pruning Agent. A higher number of threads can be used, depending on your database configuration and hardware.
- b. Specify the maximum rows that can be deleted in a single pruning transaction. Any positive integer is valid. The default value is 1000. There is no value that indicates you want all rows deleted.




If you increase the number of threads, you might consider increasing this value if your transaction log allows for it. The effective number of rows deleted per transaction is based on this value divided by the number of worker threads.

- c. Indicate a time zone for historical data from the **Use timezone offset from** drop down list. This field indicates which time zone to use when a user specifies a time period in a query for monitoring data.
 - Select **Agent** to use the time zone (or time zones) where the monitoring agents are located.
 - Select **Warehouse** to use the time zone where the Summarization and Pruning Agent is located. If the Tivoli Data Warehouse and the Summarization and Pruning Agent are in different time zones, the **Warehouse** choice indicates the time zone of the Summarization and Pruning Agent, not the warehouse.

Skip this field if the Summarization and Pruning Agent and the monitoring agents that collect data are all in the same time zone.

- d. Specify the age of the data you want summarized in the **Aggregate hourly data older than** and **Aggregate daily data older than** fields. The default value is 1 for hourly data and 0 for daily data.
 - e. The **Maximum number of node errors to display** refers to the node error table in the Summarization and Pruning workspace. It determines the maximum number of rows that workspace is to save and display.
 - f. The **Maximum number of summarization and pruning runs to display** refers to the Summarization and Pruning Run table in the Summarization and Pruning workspace. It determines the maximum number of rows that workspace is to save and display.
Maximum number of Summarization and Pruning runs to display and Maximum number of node errors to display together determine the number of rows shown in the Summarization and Pruning overall run table and Errors table respectively. There is a minimum value of 10 for each. These equate to keywords KSY_SUMMARIZATION_UNITS and KSY_NODE_ERROR_UNITS in file KSYENV/sy.ini.
 - g. The **Database Connectivity Cache Time** determines how long after a positive check for connectivity that the result will be cached. Longer times may result in inaccurate results in the workspace; however, it saves processing time.
Database Connectivity Cache Time records the number of minutes to cache the database connectivity for reporting purposes. The minimum value is 5 minutes. This equates to keyword KSY_CACHE_MINS in file KSYENV/sy.ini.
 - h. **Batch mode** determines if data from different managed systems are used in the same database batch; this setting also improves performance.
Batch mode controls the batching method used by the Summarization and Pruning Agent. A value of Single Managed System (0) means that data should only be batched for the same system, whereas a value of Multiple Managed System (1) means that data from multiple systems can be batched together; this can lead to higher performance at potentially bigger transaction sizes. The default value is Single Managed System (0). This equates to keyword KSY_BATCH_MODE in file KSYENV/sy.ini.
 - i. Specify if you want to turn **Database compression** on or off.
 - j. Specify if you want to turn **Database Table Partitioning** on or off. When partitioning is enabled, the Summarization and Pruning Agent performs the following functions:
 - Creates new summary tables as partitioned tables.
 - Adds partitions to existing partitioned tables (including the detail tables created by the warehouse proxy) based on your configuration.
 - Prunes data from partitioned tables by dropping the appropriate partitions.

For more information, see “Specifying range partitioned tables for the Summarization and Pruning Agent” on page 533.

To change these values, you can either use the Summarization and Pruning configuration window's **Additional settings** tab or update these parameters directly in file KSYENV/sy.ini.
14. Save your settings and close the window. Click **Save** to save your settings. On Windows, click **Close** to close the configuration window.
 -  Click **Save** and then click **Close**.
 -   Click **Save** and then click **Cancel**.

Configuring the Warehouse Proxy Agent on Linux or UNIX: command-line procedure

Complete the steps in this section to configure the Warehouse Proxy Agent from the command-line on Linux or UNIX.

Procedure

1. Log on to the computer where the Warehouse Proxy Agent is installed.
2. At the command-line change to the *ITMinstall_dir*/bin directory, where *ITMinstall_dir* is the directory where you installed the product.
3. Run the following command to start configuring the Warehouse Proxy Agent:

```
./itmcmd config -A hd
```

where *hd* is the product code for the Warehouse Proxy Agent.

Results

Here is a sample of Warehouse Proxy Agent configuration from the command-line. DB2 is the database used in this example, for the other supported database platforms the responses will be different:

```
itmcmd config -A hd
```

Database Type

Database [1=DB2, 2=Oracle, 3=Microsoft SQL Server] (default is: 1):

Agent Parameters :

Fully qualified paths to JDBC JAR files (comma separated)

JDBC JARs List (default is:): /data/jdbc/db2jcc.jar,/data/jdbc/db2jcc_license_cu.jar

The Warehouse JDBC URL

JDBC URL (default is: jdbc:db2://localhost:50000/WAREHOUS):

The Warehouse JDBC Driver

JDBC Driver (default is: com.ibm.db2.jcc.DB2Driver):

The Warehouse database username

Username (default is: ITMUSER):

The Warehouse database user password

Enter Password (default is:):

Re-type : Password (default is:):

Space or comma separated list of Tivoli Enterprise Monitoring Server instances served by this Warehouse Proxy agent.

*ANY can be specified if this Warehouse Proxy agent will export data of any agents connected to any TEMS. If the list is left blank, this Warehouse Proxy agent will be the default

Warehouse proxy agent.

Warehouse TEMS List (default is:): REMOTE_ITMTDWP12

Batch Database Operations

Use Batch [1=TRUE, 2=FALSE] (default is: 1):

Database Compression option

Database Compression [1=TRUE, 2=FALSE] (default is: 2):

Enable the compression of historical data from Z/OS sources before upload to the

Warehouse Proxy Server

Warehouse Compression for Z/OS Sources [1=TRUE, 2=FALSE] (default is: 2):

Enable the compression of historical data from distributed sources before upload to the

Warehouse Proxy Server

Warehouse Compression for Distributed Sources [1=TRUE, 2=FALSE] (default is: 1):

Will this agent connect to a TEMS? [1=YES, 2=NO] (Default is: 1): 1

TEMS Host Name (Default is: itmtdwp18):

Configuring the Summarization and Pruning Agent on Linux or UNIX: command-line procedure

Complete the steps in this section to configure the Summarization and Pruning Agent from the command-line on Linux or UNIX.

Procedure

1. Log on to the computer where the Summarization and Pruning Agent is installed.
2. At the command-line change to the *ITMinstall_dir*/bin directory, where *ITMinstall_dir* is the directory where you installed the product.

3. Run the following command to start configuring the Summarization and Pruning Agent:

```
./itmcmd config -A sy
```

where sy is the product code for the Summarization and Pruning Agent.

Results

Here is a sample of Summarization and Pruning Agent configuration from the command-line. DB2 is the database used in this example, for the other supported database platforms the responses will be different:

```
itmcmd config -A sy
```

Choose the database type

Database Type

Database [1=DB2, 2=Oracle, 3=Microsoft SQL Server] (default is: 1):

Sources Details

Fully qualified paths to JDBC JAR files (comma separated)

JDBC JARs List (default is: /data/jdbc/db2jcc.jar, /data/jdbc/db2jcc_license_cu.jar):

The Warehouse JDBC URL

JDBC URL (default is: jdbc:db2://localhost:50000/WAREHOUS):

The Warehouse JDBC Driver

JDBC Driver (default is: com.ibm.db2.jcc.DB2Driver):

The Warehouse user

Warehouse user (default is: itmuser):

The Warehouse password

Enter Warehouse password (default is: *):

Re-type : Warehouse password (default is: *):

The TEPS hostname

T EPS Server Host (default is: localhost):

The TEPS port (default 1920)

TEPS Server Port (default is: 1920):

Scheduling Details:

If fixed scheduling is in use

Fixed Schedule [1=No, 2=Yes] (default is: 2):

The number of days between runs (default is 1)

Every N days (default is: 1):

The fixed hour to run (valid values are 0-12, default is 2)

Hour to run (default is: 02):

The fixed minute to run (default is 0)

Minute to run (default is: 00):

AM or PM

AM/PM [1=AM, 2=PM] (default is: 1):

Minutes between flexible runs

Every N minutes (default is: 60):

Exception times in HH:MM-HH:MM format (24 hour clock), comma separated when flexible scheduling shouldn't run

Blackout (default is:):

Log Settings Details

Specify whether the WAREHOUSELOG table will be pruned. Format is nnn.unit where nnn is the number of units and unit is day, month or year. Specify blank to not prune the table.

Prune WAREHOUSELOG (default is:):

Specify whether the WAREHOUSEAGGREGLOG table will be pruned. Format is nnn.unit where nnn is the number of units and unit is day, month or year. Specify blank to not prune the table.

Prune WAREHOUSEAGGREGLOG (default is:):

Additional Settings Details

The number of worker threads to be used

Number of worker threads (default is: 2):

The maximum number of rows per transaction (effective size is this value divided by number of worker threads)

Maximum rows per database transaction (default is: 1000):

Which timezone to use when aggregating the data: agent or warehouse (default is agent)

Use timezone offset from [1=Agent, 2=Warehouse] (default is: 1):

The minimum age for hourly data to be aggregated (default is 1)

Aggregate hourly data older than (default is: 1):

The minimum age for daily data to be aggregated (default is 0)

Aggregate daily data older than (default is: 0):

The number of errors to keep in memory (default is 10)

Maximum number of node errors to display (default is: 10):

The number of summarization runs to keep in memory (default is 10)

Maximum number of Summarization and Pruning runs to display (default is: 10):

The number of minutes to cache the database status (default is 10)

Database Connectivity Cache Time (minutes) (default is: 10):

The type of batching to be used (default is single system)

Batch mode [1=Single System, 2=Multiple System] (default is: 1):

Enable database compression, if supported (default is no)

Database compression [1=No, 2=Yes] (default is: 1):

Will this agent connect to a TEMS? [1=YES, 2=NO] (Default is: 1):

TEMS Host Name (Default is: itmtdwp18):

Starting the Summarization and Pruning Agent

This section contains instructions for starting the Summarization and Pruning Agent.

- To start the Summarization and Pruning Agent from the Manage Tivoli Enterprise Monitoring Services window, right-click **Summarization and Pruning** and select **Start**.
- (*Linux or UNIX only*) To start the Summarization and Pruning Agent from the command-line, run the following command from the bin directory of the IBM Tivoli Monitoring installation directory. The default installation directory is /opt/IBM/ITM.

```
./itmcmd agent start sy
```

where sy is the product code for the Summarization and Pruning Agent.

Installing and configuring multiple Warehouse Proxy Agents

IBM Tivoli Monitoring supports multiple Warehouse Proxies within a single hub monitoring server environment. The provision for multiple warehouse proxies provides for greater scalability and performance in historical data collection.

More importantly, use multiple proxies enhance reliability via their failover mechanism: if a Warehouse Proxy is unavailable, data can be inserted into the warehouse by a different Warehouse Proxy Agent (if the agents are configured properly for failover).

The support for multiple Warehouse Proxy Agents has the following important features:

- All Warehouse Proxy Agents within a single hub monitoring server environment export data to a single Tivoli Data Warehouse.
- Each Warehouse Proxy Agent is associated with a subset of monitoring server instances that you specify when you configure the proxy agent. Each Warehouse Proxy exports data only for monitoring agents that report to one of the monitoring servers (hub or remote) on the specified list. Alternatively, a given Warehouse Proxy Agent can be configured to service *any* monitoring server. This is important for backward compatibility and for configuring failover.

The following sequence of events explains how the monitoring agents, which collect the data for historical reports, know which Warehouse Proxy Agent to use:

1. When a Warehouse Proxy Agent starts, it registers with the Global Location Broker on the hub monitoring server, sending it the list of monitoring servers that it is configured to serve, or an indication that the Warehouse Proxy Agent can service any monitoring server. This registration process is repeated every hour.
2. Each monitoring server queries the Global Location Broker at regular intervals to determine which Warehouse Proxy it is associated with. The monitoring server then sends the address of this Warehouse Proxy to all of its child monitoring agents to use during historical data exports. You can change the default query interval of 60 minutes to some other value.

When a Warehouse Proxy Agent registers with the Global Location Broker, it is registered as the default proxy agent if no other proxy agent is already configured as the default. When a monitoring server queries the Global Location Broker for its associated Warehouse Proxy, the default proxy agent is used if that monitoring server is not on the list of servers for any proxy agent.

Setting a permanent socket address for a proxy agent

In some network environments, use of the Global Location Broker default registration algorithm is not supportable and a specific monitoring server variable must be used.

This applies also when monitoring agents configure their connections as ephemeral (using `KDC_FAMILIES=ip.pipe use:y ephemeral:y` for instance).

KPX_WAREHOUSE_LOCATION is the variable that allows a fixed warehouse route to be delivered to connected agents. KPX_WAREHOUSE_LOCATION is an optional list of fully qualified, semicolon delimited network names and should be added in the monitoring server environment file located in different place depending on the platform:

- **Windows** `install_dir\CMS\KBBENV`
- **Linux** **UNIX** `install_dir/config/kbbenv.ini`

KPX_WAREHOUSE_LOCATION is used instead of the routing string currently derived by the monitoring server from the Location Broker monitoring. If the variable KPX_WAREHOUSE_LOCATION is set on the hub monitoring server, then the Warehouse Proxy Agent is registered in the Global Location Broker. If the variable is set on the remote monitoring server, then the Warehouse Proxy Agent is registered in the Local Location Broker. This variable has a maximum length of 200 bytes. The syntax is:

`KPX_WAREHOUSE_LOCATION= family_protocol:network_address[port number];`

Note: For Tivoli Enterprise Monitoring Servers running on UNIX and Linux platforms only, you must reconfigure the Tivoli Enterprise Monitoring Server after adding KPX_WAREHOUSE_LOCATION to the `kbbenv.ini` file. This is required so that the new setting is propagated to the KBBENV file.

Verifying the configuration

Use the procedures in this section to verify the configuration and to determine which Warehouse Proxy a particular monitoring server uses for its agents.

Use the following trace settings to verify the configuration:

- To verify that a Warehouse Proxy is registering with the hub monitoring server and placing the correct entries into the Global Location Broker:
 1. Open the environment file for the proxy agent:
 - (Windows) `ITMinstall_dir\TMAITM6\KHDENV`
 - (Linux or AIX) `ITMinstall_dir/config/hd.ini`
 where `ITMinstall_dir` is the directory where you installed the product.
 2. Add the following entry to the KBB_RAS1 trace setting:


```
KBB_RAS1=ERROR(UNIT:khdxrpcr STATE)
```

This setting prints the value of KHD_WAREHOUSE_TEMS_LIST and shows any errors associated with its components.

- To determine which Warehouse Proxy a particular monitoring server uses for its agents:

1. Open the environment file for the monitoring server:

- (Windows) *ITMinstall_dir*\CMS\KBBENV
- (Linux or UNIX) *ITMinstall_dir*/config/hostname_ms_TEMSid

where *ITMinstall_dir* is the directory where you installed the product.

2. Add the following entry to the KBB_RAS1 trace setting:

```
KBB_RAS1=ERROR(UNIT:kpxrwhpx STATE)
```

This setting prints entries in the RAS log of the monitoring server when a registration change occurs. The entry specifies the name and address of the new Warehouse Proxy Agent that the monitoring server is using.

Running the warehouse agents autonomously

Both the Warehouse Proxy Agent and Warehouse Summarization and Pruning Agent can run in autonomous ("unmanaged") mode.

If a Warehouse Proxy Agent is running in connected ("managed") mode, it makes its location available to the monitoring servers and application monitoring agents by registering the location with the Global Location Broker on the hub monitoring server, along with a list of monitoring servers that it supports. Remote monitoring servers and the agents that connect to them obtain the network addresses of the proxy agents to which they report from the hub. However, if the Warehouse Proxy Agent is running in autonomous ("unmanaged") mode, it does not register its location with the hub monitoring server; instead, monitoring agents are configured to obtain the location of their Warehouse Proxy Agents from their local configuration file or from a central configuration server facility. This allows the Warehouse Proxy Agent to support monitoring agents that are running autonomously (that is, without a connection to a monitoring server).

If the Summarization and Pruning Agent is running in connected mode, it obtains the information it needs about the collected historical data from the Tivoli Enterprise Portal Server. If the Summarization and Pruning Agent is configured to run in autonomous mode, it obtains the information about attribute data directly from attribute support files for the monitoring agents, and the summarization and pruning settings from the WAREHOUSESUMPRUNE table in the warehouse database. The location of the application support files is specified in the Summarization and Pruning Agent configuration file (see "Configuring a Summarization and Pruning Agent to run autonomously" on page 659). See "Configuring summarization and pruning without the Tivoli Enterprise Portal Server" on page 660 for more information on the WAREHOUSESUMPRUNE table.

Configuring a Warehouse Proxy Agent to run autonomously

For a Warehouse Proxy Agent to run in autonomous mode, it must be configured to run without registering its location with the hub monitoring server.

Every monitoring agent that sends historical data to the Warehouse Proxy Agent must be configured with its location.

To configure the Warehouse Proxy Agent to run without registering, complete the following steps:

1. Install and configure one or more Warehouse Proxy Agents following the procedures in the appropriate chapter:
 - Chapter 23, "Tivoli Data Warehouse solution using DB2 for Linux, UNIX, and Windows," on page 539
 - Chapter 25, "Tivoli Data Warehouse solution using Microsoft SQL Server," on page 595

- Chapter 26, “Tivoli Data Warehouse solution using Oracle,” on page 619

Note: The installation procedure for Windows operating systems includes steps for configuring the connection between the proxy agent and the hub Tivoli Enterprise Monitoring Server. On Linux or UNIX operating systems, this step is performed in a separate configuration procedure. On Windows operating systems, if you want to run the Warehouse Proxy Agent without a connection to the hub, accept the defaults for the connection information, but specify a nonvalid name for the monitoring server. On UNIX and Linux operating systems, check *No TEMS* on the **TEMS Connection** tab of the configuration window.

2. Add the following variable to the Warehouse Proxy Agent environment file (*install_dir\TMAITM6\KHDENV* on Windows operating systems; *install_dir/config/hd.ini* on UNIX and Linux operating systems):

```
KHD_REGWITHGLB=N
```

3. Configure the Warehouse Proxy Agent to use the same IP port number as you chose for the various autonomous agents that will be sending historical data to it; see the *IBM Tivoli Monitoring Administrator's Guide* for details.

4. Restart the agent.

5. Restart the agent's Tivoli Enterprise Monitoring Server, if necessary.

If the Warehouse Proxy Agent you reconfigured to run autonomously has previously connected to either a hub monitoring server or a remote monitoring server, the agent has already registered with the monitoring server to which it connected. To clear this registration information now that the agent is running autonomously, recycle the monitoring server. If the monitoring server is a remote monitoring server, also recycle the hub monitoring server to which it connects.

To configure a monitoring agent with the location of the Warehouse Proxy Agent or agents to which it should export historical data, complete the following steps:

1. Install the agent following “Installing monitoring agents” on page 261 or the documentation for the agent.

2. Open the monitoring agent environment file in a text editor:

- **Windows** *install_dir\TMAITM6\kpcenv*
- **Linux** **UNIX** *install_dir/config/pc.ini*
- **z/OS** *&chilev.&rte.RKANPARU(KPCENV)*

where *pc* is the two-character product code for the monitoring agent (see Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883).

3. Add the following variable to the file:

```
KHD_WAREHOUSE_LOCATION=family protocol.#network address[port number]
```

The value of the variable can be a semicolon-delimited list of network addresses. For example:

```
KHD_WAREHOUSE_LOCATION=ip.pipe:SYS2-XP[63358];ip:SYS2-XP[63358]
```

4. Restart the agent.

Configuring a Summarization and Pruning Agent to run autonomously

For a Summarization and Pruning Agent to run without connecting to a Tivoli Enterprise Portal Server, it must be configured to run autonomously and it must be able to locate the application support files for the monitoring agents.

About this task

To configure a Summarization and Pruning Agent to run in autonomous mode, complete the following steps:

Procedure

1. If you have not already installed done so, install a Tivoli Enterprise Portal Server, and add application support for all types of monitoring agents that will be collecting historical data.
2. Install and configure Summarization and Pruning Agent following the instructions in the appropriate chapter:
 - Chapter 23, “Tivoli Data Warehouse solution using DB2 for Linux, UNIX, and Windows,” on page 539
 - Chapter 25, “Tivoli Data Warehouse solution using Microsoft SQL Server,” on page 595
 - Chapter 26, “Tivoli Data Warehouse solution using Oracle,” on page 619
3. If the Summarization and Pruning Agent is not installed on the same machine as the portal server, copy the required application support files to the machine on which the Summarization and Pruning agent resides.

These files are named `dockpc`, where `pc` is the two-letter product code for the monitoring agent (see Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883).

Windows The files are located in `install_dir\cnps` directory; on **Linux** **UNIX** The files are located in the `install_dir/arch/cq/data` directory.

By default, the Summarization and Pruning Agent looks for the ODI files in the `install_dir\TMAITM6` directory (on Windows) or the `install_dir/arch/cq/data` directory (on UNIX and Linux). If you do not create this directory and copy the files to it, you must add the **KSY_AUTONOMOUS_ODI_DIR** variable to the Summarization and Pruning Agent environment file and specify the alternative location.

Note: There is no need to copy the `dockcj` file; it is not used when reconfiguring the Summarization and Pruning Agent. If you do copy this file, the following error will occur and can be ignored.

Validation failed: Column name exceeds 10 characters: ACKNOWLEDGED.
ODI File contents not loaded: `/install_dir/dockcj`

4. On the machine where the Summarization and Pruning Agent is installed, open its environment file in a text editor:
 - **Windows** `install_dir\TMAITM6\KSYENV`
 - **Linux** **UNIX** `install_dir/config/sy.ini`
5. Edit the following variables:
 - To enable the Summarization and Pruning Agent to run without connecting to the Tivoli Enterprise Portal Server, set **KSY_AUTONOMOUS=Y**.
 - If you did not install the application support files in the default directory (see step 3), set **KSY_AUTONOMOUS_ODI_DIR=<alternative location of application support files>**.
6. Restart the Summarization and Pruning Agent agent. The **WAREHOUSESUMPRUNE** table is automatically created when the Summarization and Pruning Agent is started.
7. If you are upgrading from a previous version and already have summarization and pruning settings stored in the Tivoli Enterprise Portal Server database, restart the Tivoli Enterprise Portal Server.

The first time the portal server is started after the **WAREHOUSESUMPRUNE** table has been created, any previously existing data collection and summarization and pruning configuration settings are migrated to the **WAREHOUSESUMPRUNE** table in the warehouse data base. Subsequently, any settings configured using the portal server are stored directly in the warehouse.

Configuring summarization and pruning without the Tivoli Enterprise Portal Server

You can configure historical data collection and summarization and pruning using the Tivoli Enterprise Portal, or you can configure them directly in the warehouse database **WAREHOUSESUMPRUNE** table using SQL commands.

Table 124 on page 661 contains descriptions of the columns in the **WAREHOUSESUMPRUNE** table. Insert one row for each attribute group for which you want to collect historical data, along with the values for

any summarization and pruning settings. You do not need to set defaults for unused options; they are built into the table design. Varchar values must be enclosed in single quotes (' ')

Table 124. Descriptions of the columns in the WAREHOUSESUMPRUNE control settings table

Name	Type	Description
TABNAME	VARCHAR (40) NOT NULL PRIMARY KEY	The short table name. In the application support file, this is the value of TABLE. Review the application support file associated with each agent for the TABLE names.
YEARSUM	VARCHAR (8) DEFAULT '-16823'	Yearly Summarization on (-16822); off (-16823)
QUARTSUM	VARCHAR (8) DEFAULT '-16823'	Quarterly Summarization on (-16822); off (-16823)
MONSUM	VARCHAR (8) DEFAULT '-16823'	Monthly Summarization on (-16822); off (-16823)
WEEKSUM	VARCHAR (8) DEFAULT '-16823'	Weekly Summarization on (-16822); off (-16823)
DAYSUM	VARCHAR (8) DEFAULT '-16823'	Hourly Summarization on (-16822); off (-16823)
HOURSUM	VARCHAR (8) DEFAULT '-16823'	Daily Summarization on (-16822); off (-16823)
PYEAR	VARCHAR (8) DEFAULT '-16838'	Yearly Pruning on (-16837); off (-16838)
PYEARINT	SMALLINT DEFAULT 1	Number of units (years, months, or days)
PYEARUNIT	VARCHAR (8) DEFAULT '-16834'	Type of Unit. Years (-16834), Months (-16835), Days (-16836)
PQUART	VARCHAR (8) DEFAULT '-16838'	Quarterly Pruning on (-16837); off (-16838)
PMON	VARCHAR (8) DEFAULT '-16838'	Monthly Pruning on (-16837); off (-16838)
PMONINT	SMALLINT DEFAULT 1	Number of units (years, months, or days)
PMONUNIT	VARCHAR (8) DEFAULT '-16835'	Type of Unit. Years (-16834), Months (-16835), Days (-16836)
PWEEK	VARCHAR (8) DEFAULT '-16838'	Weekly Pruning on (-16837); off (-16838)
PWEEKINT	SMALLINT DEFAULT 1	Number of units (years, months or days)
PWEEKUNIT	VARCHAR (8) DEFAULT '-16835'	Type of Unit. Years (-16834), Months(-16835), Days(-16836)
PDAY	VARCHAR (8) DEFAULT '-16838'	Daily Pruning on (-16837); off (-16838)
PDAYINT	SMALLINT DEFAULT 1	Number of units (years, months, or days)
PDAYUNIT	VARCHAR (8) DEFAULT '-16835'	Type of Unit. Years (-16834), Months (-16835), Days (-16836)
PHOUR	VARCHAR (8) DEFAULT '-16838'	Hourly Pruning on (-16837); off (-16838)
PHOURINT	SMALLINT DEFAULT 1	Number of units (years, months, or days)
PHOURUNIT	VARCHAR (8) DEFAULT '-16836'	Type of Unit. Years (-16834), Months (-16835), Days (-16836)
PRAW	VARCHAR (8) DEFAULT '-16838'	Detailed Pruning on (-16837); off (-16838)

Table 124. Descriptions of the columns in the WAREHOUSESUMPRUNE control settings table (continued)

Name	Type	Description
PRAWINT	SMALLINT DEFAULT 1	Number of units (years, months, or days)
PRAWUNIT	VARCHAR (8) DEFAULT '-16836'	Type of Unit. Years (-16834), Months (-16835), Days (-16836)

Examples:

This section contains examples of basic collection and summarization and pruning configuration.

Configuration and daily/hourly summarization

Collection is configured, and daily and hourly summarizations are set. No pruning has been specified. Use the SQL INSERT command.

Required:

- TABNAME= *Table Code*
- DAYSUM= -16822 (summarize daily)
- HOURSUM=-16822 (summarize hourly)

```
INSERT INTO WAREHOUSESUMPRUNE (TABNAME,DAYSUM,HOURSUM) VALUES ('WTMEMORY',' -16822',' -16822');
```

Configuration, daily/hourly summarization, and daily/hourly pruning

Collection is configured, and daily and hourly summarizations are set. Pruning is specified for daily 3-month intervals and hourly 2-day intervals. Use the SQL INSERT command.

Required:

- TABNAME=*Table Code*
- DAYSUM=-16822 (summarize daily)
- HOURSUM= -16822 (summarize hourly on)
- PDAY= -16837 (prune daily)
- PDAYINT=3 (prune interval)
- PDAYUNIT= -16835 (prune monthly)
- PHOUR = -16837 (prune hourly on)
- PHOURINT= 2 (prune interval)
- PHOURUNIT= -16836 (prune daily)
- PRAW = -16837 (prune detailed on)
- PRAWINT= 1 (prune interval)
- PRAWUNIT= -16836 (prune daily)

```
INSERT INTO WAREHOUSESUMPRUNE
(TABNAME,DAYSUM,HOURSUM,PDAY,PDAYINT,PDAYUNIT,PHOUR,PHOURINT,PHOURUNIT,PRAW,PRAWINT,PRAWUNIT)
VALUES ('WTMEMORY',' -16822',' -16822',' -16837',3,' -16835',2,' -16836',' -16837',1,' -16836');
```

Remove daily summarization and pruning

In this example, collection has been configured, and daily and hourly summarizations have been set. You want to remove daily summarization and pruning. Use the SQL UPDATE command.

Required:

- TABNAME= *Table Code*
- DAYSUM= -16823 (summarize off)
- PDAY= -16838 (prune daily off)
- PDAYINT= 1 (prune interval to default)
- PDAYUNIT= -16836 (prune daily)

```
UPDATE WAREHOUSESUMPRUNE SET DAYSUM='-16823', SET PDAY='-16838', SET PDAYINT=1,  
SET PDAYUNIT='-16836' WHERE TABNAME='WTMEMORY';
```

Remove collection entirely

Collection has been configured, and daily and hourly summarizations have been set. Set the collection to unconfigured by removing the entire row from the table, using the SQL DELETE command.

Required: TABNAME= *Table Code*

```
DELETE FROM WAREHOUSESUMPRUNE WHERE TABNAME='WTMEMORY';
```

Testing the connection between the portal server and the Tivoli Data Warehouse

To test the connection between the Tivoli Enterprise Portal Server and the Tivoli Data Warehouse database, use the Tivoli Enterprise Portal to create a custom SQL query to the warehouse database and display the results. The procedure described in this section creates an SQL query to the WAREHOUSELOG table, a status table that is created in the warehouse database when the Warehouse Proxy Agent starts.

See “WAREHOUSELOG and WAREHOUSEAGGREGLOG tables” on page 689 for a description of this table.

Before you begin, make sure that the following components are installed and started:

- The Tivoli Enterprise Portal Server
- The Warehouse Proxy Agent
- The Tivoli Data Warehouse RDBMS server



The test procedure consists of the following steps:

1. Create a custom SQL query to the WAREHOUSELOG table.
2. Create a new workspace to display the results.
3. Assign the query to the workspace.

1. Create the query

The first step of the test procedure is to create a custom SQL query to the WAREHOUSELOG table.

Procedure

1. Start the Tivoli Enterprise Portal using the desktop or browser client. See “Starting the Tivoli Enterprise Portal client” on page 366 for instructions.
Log in as the system administrator. (The default system administrator ID is **sysadmin**.)
2. Click  **Query Editor** on the Tivoli Enterprise Portal tool bar.
3. In the Query Editor window, click  **Create New Query**.
The Create Query window is displayed.





Figure 143. Create Query window

4. Enter a name and description for the query. For this test, enter WAREHOUSELOG for the query name.
5. From the **Category** drop-down list, select the folder where you want the WAREHOUSELOG query to appear in the Query tree.
For example, select a folder name that corresponds to the operating system (Windows OS or UNIX OS) where the Tivoli Data Warehouse is installed.
6. Select the name of the data source for the Tivoli Data Warehouse in the **Data Sources** list.
7. Click **OK**.
The WAREHOUSELOG query appears in the Query tree in the Custom_SQL folder under the category that you selected. The **Specification** tab opens with a **Custom SQL** text box for you to enter an SQL command.
8. Enter the following SQL statement in the **Custom SQL** text box to select all columns of the WAREHOUSELOG table:
select * from WAREHOUSELOG
9. Click **OK** to save the query and close the Query Editor window.

2. Create a workspace

The second step of the test procedure is to create a workspace.

Procedure


1. If the Enterprise Status workspace is not open, click the  Enterprise Navigator item.
2. Select **Save Workspace As** from the File menu.
3. Enter the name WAREHOUSELOG for the new workspace.
4. Click **OK**.

You now have a duplicate of the Enterprise Status workspace named WAREHOUSELOG that you can modify.

3. Assign the query to the workspace

The third step of the test procedure is to assign the query to the workspace.

Procedure

1. In the WAREHOUSELOG workspace, click  **Table** on the tool bar, release the mouse button; then click inside one of the workspace views (for example, the Situation Event Console view).
2. Click **Yes** on the message asking if you want to assign the query now.
3. In the Query tab, select **Click here to assign query**.
The Query editor opens with a list of all product queries.
4. In the list of queries, expand the category folder you specified when you created the query (for example Windows OS or UNIX OS), and then expand the Custom_SQL folder.
5. Select **WAREHOUSELOG** from the list of queries in the Custom_SQL folder.

The query definition opens in the right frame.




Figure 144. Assigning the WAREHOUSELOG query to a workspace

6. Click **OK** to select the WAREHOUSELOG query for the table and close the Query editor.
7. Click **OK** to close the **Properties** editor.

The columns of the WAREHOUSELOG table are displayed within the table view in the WAREHOUSELOG workspace. The content of the table is also displayed if the table is not empty.

Tuning the performance of the Warehouse Proxy

This section contains the environment variables that control how the Warehouse Proxy works.

You can set the following variables in the KHDENV file:

- “Database initialization” on page 667
- “Work queue” on page 667
- “Connection pool” on page 667
- “RPC threads and export requests” on page 668
- “Timeout values” on page 668

Database initialization

This section contains the tests and environment variables used during database initialization.

When the Warehouse Proxy starts, the following tests are done:

- Checks that the Warehouse Proxy can connect to the database.
- If the database is Oracle or DB2 for Linux, UNIX, and Windows, checks that the encoding is set to UTF8.
- If the database is DB2 for Linux, UNIX, and Windows, checks that a bufferpool of page size 8KB is created. If it is not, one is created, along with three new table spaces that use the 8KB bufferpool. The bufferpool is called "ITMBUF8K" and the table spaces are named "ITMREG8K," "ITMSYS8K," and "ITMBUF8K."
- Creates a database cache that contains a list of all the tables and columns that exist in the database.

If any of these tests fail, a message is written to the log file and messages are displayed in the Event Viewer.

These tests are repeated every 10 minutes.

You can change this default start up behavior by changing the following environment variables:

KHD_CNX_WAIT_ENABLE

Enables the Warehouse Proxy to wait in between attempts to connect to the database. By default, this variable is set to Y. If you do not want the Warehouse Proxy to wait, change the variable to N. However, this can generate a large log file if the connection to the database fails with each attempt.

KHD_CNX_WAIT

Defines the amount of time, in minutes, that the Warehouse Proxy waits between attempts to connect to the database. The default value is 10 minutes.

Work queue

The work queue consists of a single queue instance and a configurable number of worker threads that run work placed on it.

There are two primary configuration parameters that you can set. You can set these parameters in the KHDENV file on Windows or the hd.ini file on Linux and UNIX before starting the Warehouse Proxy.

KHD_QUEUE_LENGTH

The length of the KHD work queue. This is an integer that identifies the maximum number of export work requests that can be placed on the work queue before the work queue begins rejecting requests. The default value is 1000. Setting this value to 0 means that the work queue has no limit.

KHD_EXPORT_THREADS

The number of worker threads exporting data to the database. The default value is 10.

Connection pool

The Warehouse Proxy uses several pre-initialized ODBC connections to access the target database. The use of these ODBC connection objects is synchronized through a single connection pool. The connection pool is initialized when the Warehouse Proxy starts.

You can configure the number of connections in the pool by defining the following environment variable in the KHDENV file on Windows or the hd.ini file on Linux and UNIX:

- KHD_CNX_POOL_SIZE: The total number of pre-initialized ODBC connection objects available to the work queue export threads. The default value is 10.

All export worker threads request connections from the connection pool and must obtain a connection before the work of exporting data can continue.

You only see the connections established when a request is active. It is important to set the number of worker threads to greater or equal to the number of ODBC connections. To do this, set `KHD_KHD_EXPORT_THREADS >= KHD_CNX_POOL_SIZE`.

RPC threads and export requests

In addition to the configurable environment variables discussed previously, the standard agent framework provides some control over the scalability and performance profile for the Warehouse Proxy.

When the Warehouse Proxy starts, it initializes RPC and registers a group of function pointers that respond to incoming RPC calls.

Use the `CTIRA_NCSLISTEN` variable to set the number of RPC threads.

Timeout values

You can set two environment variables to control the timeout value. One variable is set on the application agent, the other on the Warehouse Proxy.

KHD_STATUSTIMEOUT

The time, in seconds, to wait for a status from the Warehouse Proxy before sending an export request again. The default value is 900 seconds, or 15 minutes.

KHD_SRV_STATUSTIMEOUT

The timeout value, in seconds, for the work queue to perform work. The default value is 600 seconds, or 10 minutes.

Export requests are rejected by the Warehouse Proxy are the following four reasons:

- The time between when an export request is sent to the work queue and when it is extracted from the queue exceeds the timeout. If you have tracing for the Warehouse Proxy set to ERROR, an error similar to the following is logged in the Warehouse Proxy log file:

REJECTED: The export for the originnode *OriginNodeName*, the application *applicationName* and the table *tableName* has been rejected for timeout reason in stage END_QUEUE.

- The time between when an export request is sent to the work queue and when the work queue starts to do existence checking in the database exceeds the timeout. If you have tracing for the Warehouse Proxy set to ERROR, an error similar to the following is logged in the Warehouse Proxy log file and the WAREHOUSELOG table:

Sample data rejected for timeout reason at stage START_EXPORT

- The time between when an export request is sent to the work queue and when the work queue fetches all the rows in the sample exceeds the timeout. If you have tracing for the Warehouse Proxy set to ERROR, a message similar to the following is logged in the Warehouse Proxy log file and the WAREHOUSELOG table:

Sample data rejected for timeout reason at stage START_SAMPLE

- The time between when an export request is sent to the work queue and when the work queue commits the rows in the database exceeds the timeout. If you have tracing for the Warehouse Proxy set to ERROR, a message similar to the following is logged in the Warehouse Proxy log file and the WAREHOUSELOG table:

Sample data rejected for timeout reason at stage COMMIT

The `KHD_SRV_STATUSTIMEOUT` variable should be set less than `KHD_STATUSTIMEOUT` by at least 60 seconds. Do not change these values unless there is a problem in the environment.

Configuring the Warehouse Proxy Agent and Summarization and Pruning Agent using the new graphical interface

The database administrator must create an empty database that will serve as the Data Warehouse database and provide the database name, username and password to the IBM Tivoli Monitoring administrator. The IBM Tivoli Monitoring administrator creates the ODBC or JDBC Data Source using the credentials provided. Then, the IBM Tivoli Monitoring administrator must configure the Warehouse Proxy Agent and the Summarization and Pruning Agent to use the ODBC or JDBC Data Source. The Warehouse Proxy Agent uses ODBC on Windows systems. On Unix and Linux, the Warehouse Proxy Agent uses JDBC. The Summarization and Pruning Agent always uses JDBC.

You can confirm a successful configuration when the Warehouse Proxy Agent and Summarization and Pruning Agent can connect to the Data Warehouse database. To verify this, login to the Tivoli Enterprise Portal Desktop and check the state of the Data Warehouse database connectivity for both the agents. These should be displayed as green.

You need to complete the following steps for this configuration option:

1. During installation, when the Configuration Defaults for Connecting to a Tivoli Enterprise Monitoring Server window is displayed, you will see the user interface for the Warehouse Proxy Agent. Select the database type and click **Next**.
2. Enter the database connection information and adjust any other configuration parameters, then click **OK**.

Installing and configuring the Warehouse Proxy Agent and Summarization and Pruning Agent in silent mode, using the silent response file

A silent installation runs on its own and does not require you to provide input to dialog boxes or monitor the installation. This configuration option allows for an unattended and automatic installation and configuration of the Warehouse Proxy Agent and Summarization and Pruning Agent as well as the remote configuration of these two agents. Configuration of the Warehouse Proxy Agent and Summarization and Pruning Agent takes place during the setup of the IBM Tivoli Monitoring infrastructure, or manually later.

Before you begin

The database administrator must create the following before the silent install is run:

- An empty database that will serve as the Data Warehouse database. The default DB2 database is WAREHOUS.
- The username and password that will be used with the database default ITMUser.
- The ODBC connection that will be used to connect the database to the Warehouse Proxy and Summarization and Pruning Agents default ITM Warehouse. The ODBC datasource connection must be created in 32-bit and 64-bit mode if the 64-bit Warehouse Proxy Agent is used to connect to the Data Warehouse database.

After these items are created, the database administrator must provide the database name, username and password, and ODBC connection information to the IBM Tivoli Monitoring administrator. The IBM Tivoli Monitoring administrator takes the information provided by the database administrator and updates all the configuration information in the silent response file that drives the installation and configuration process.

About this task

You must complete the following steps for this configuration option:

Procedure

1. Fill in the silent response file (`silent_server.txt`), providing the configuration information for the Warehouse Proxy Agent and the Summarization and Pruning Agent.

You can also have response files created automatically on Windows. From the Manage Tivoli Enterprise Monitoring Services screen, right-click the agent whose configuration information you want saved, and select the **Advanced** → **Utilities** → **Generate Response Files** option from the pop-up menu.

2. Start the installation in silent mode, providing the silent response file as a parameter.
 - a. Start a DOS Command Shell.
 - b. From the shell, cd to the directory containing this installation (where `setup.exe` and `setup.isn` reside).
 - c. Invoke setup as follows:

```
start /wait setup /z"/sfC:\temp\SILENT_SERVER.txt" /s /f2"C:\temp\silent_setup.log"
```

You must specify the parameters in the same order listed above.

Where:

/z"/sf

Specifies the name of the installation driver you have customized for your site. This is a required parameter. This file must exist.

/s Specifies that this is a silent install. This causes no responses to be displayed during installation on the installed target workstation.

/f2

Specifies the name of the InstallShield log file. If you do not specify this parameter, the default is to create `Setup.log` in the same location as `setup.isn`. This log is the InstallShield log not the install log. The install log can be found in the install target directory, default `c:\IBM\ITM` sub-directory `InstallITM`, or on the Windows Boot drive root directory if the install aborts before the install location has been identified. In either case, the Setup program must be able to create and write to this file.

What to do next

You can confirm a successful configuration when the Warehouse Proxy Agent, Summarization and Pruning Agent, and Performance Analyzer Warehouse Agent show a connection to the Data Warehouse database. To verify this, login to the Tivoli Enterprise Portal Desktop and check the state of the Data Warehouse database connectivity for all three agents. The Summarization and Pruning Agent should also show connectivity to the Tivoli Enterprise Portal. The connections should be displayed as green.

Installing and configuring the Warehouse Proxy Agent in silent mode on UNIX, using the silent response file

A silent installation runs on its own and does not require you to provide input to dialog boxes or monitor the installation. This reduces the time spent on installation and configuration.

About this task

To install and configure the Warehouse Proxy Agent on UNIX with a silent response file, you must complete the following steps:

Procedure

1. Create the silent install response file, see “Warehouse Proxy Agent silent install response file sample” on page 674. Remember that user and password information are not encrypted in the response file.
2. Run the following command from the installation image:
`./install.sh -q -h <install_dir> -p <response_file>`

Note: You must not specify the path of the directory containing `./install.sh` as your IBM Tivoli Monitoring home directory. On certain platforms, this can cause the plug-in JAR files to overwrite themselves and become zero length files. The installation will fail as a result.

In addition, do not use the number sign (#) or exclamation mark (!) characters anywhere in your installation directory path.

3. Edit the silent config response file from the samples directory and change it according to your configuration, see “Warehouse Proxy Agent silent config response file sample” on page 674.
4. Run the silent config command to configure the Warehouse Proxy Agent:
`itmcmd config -A -h <install_dir> -p <response file> hd`

Results

Note: After running the command in step 2, the Warehouse Proxy Agent should be installed but not yet configured. The install `itm_install_output` log file should display content similar to the following:

You selected the following products:
Warehouse Proxy

```
... installing "Warehouse Proxy for AIX R5.3 (64 bit) AIX R6.1 (64 bit)";  
please wait.
```

```
=> installed "Warehouse Proxy for AIX R5.3 (64 bit) AIX R6.1 (64 bit)".  
... Initializing component Warehouse Proxy for AIX R5.3 (64 bit) AIX R6.1 (64 bit).  
... Warehouse Proxy for AIX R5.3 (64 bit) AIX R6.1 (64 bit) initialized.
```

```
... postprocessing; please wait.  
... finished postprocessing.  
Installation step complete.
```

Note: After running the command in step 4, the config trace file should show the following:

```
<date> ITMinstall.CandleConfigAgent main [LOG_INFO]  
Agent configuration completed...
```

The config file `hd.ini` should contain the exact information contained in the response file.

What to do next

After the Warehouse Proxy Agent has been installed and configured, it can be started using the command-line `itmcmd agent start hd`. You should check that the Agent started without any errors.

Installing and configuring the Summarization and Pruning Agent in silent mode on UNIX, using the silent response file

A silent installation runs on its own and does not require you to provide input to dialog boxes or monitor the installation. This reduces the time spent on installation and configuration.

About this task

To install and configure the Summarization and Pruning Agent on UNIX with a silent response file, you must complete the following steps:

Procedure

1. Create the silent install response file, see “Summarization and Pruning Agent silent install response file sample” on page 677. Remember that user and password information are not encrypted in the response file.
2. Run the following command from the installation image:
`./install.sh -q -h <install_dir> -p <response_file>`

Note: You must not specify the path of the directory containing `./install.sh` as your IBM Tivoli Monitoring home directory. On certain platforms, this can cause the plug-in JAR files to overwrite themselves and become zero length files. The installation will fail as a result.

In addition, do not use the number sign (#) or exclamation mark (!) characters anywhere in your installation directory path.

3. Edit the silent config response file from the samples directory and change it according to your configuration, see “Summarization and Pruning Agent silent config response file sample” on page 677.
4. Run the silent config command to configure the Summarization and Pruning Agent:
`itmcmd config -A -h <install_dir> -p <response file> sy`

Results

Note: After running the command in step 2, the Summarization and Pruning Agent should be installed but not yet configured. The install `itm_install_output` log file should display content similar to the following:

You selected the following products:
Summarization and Pruning Agent

```
..installing "Summarization and Pruning Agent for AIX R5.3(64 bit) AIX R6.1
(64 bit)"; please wait.
```

```
=> installed "Summarization and Pruning Agent for AIX R5.3 (64 bit)
AIX R6.1 (64 bit)".
..Initializing component Summarization and Pruning Agent for AIX R5.3 (64 bit)
AIX R6.1 (64 bit).
..Summarization and Pruning Agent for AIX R5.3 (64 bit) AIX R6.1 (64 bit)
initialized.
```

```
..postprocessing; please wait.
..finished postprocessing.
Installation step complete.
```

Note: After running the command in step 4, the config trace file should show the following:

```
<date> ITMinstall.CandleConfigAgent main [LOG_INFO]
Agent configuration completed...
```

The config file `hd.ini` should contain the exact information contained in the response file.

What to do next

After the Summarization and Pruning Agent has been installed and configured, it can be started using the command-line `itmcmd agent start sy`. You should check that the Agent started without any errors.

Upgrading the Warehouse Proxy Agent and Summarization and Pruning Agent in silent mode, using the silent response file

A silent installation runs on its own and does not require you to provide input to dialog boxes or monitor the installation. This reduces the time spent on upgrading.

About this task

To upgrade the Warehouse Proxy Agent or the Summarization and Pruning Agent in silent mode, you must complete the following steps:

Procedure

1. The Warehouse Proxy Agent or Summarization and Pruning Agent should already be installed at a lower level before upgrading.
2. All agents should be stopped before upgrading.
3. Create the silent install response file, see “Warehouse Proxy Agent silent install response file sample” on page 674 or “Summarization and Pruning Agent silent install response file sample” on page 677.
4. Run the following command from the installation image:

```
./install.sh -q -h <install_dir> -p <response_file>
```

Note: You must not specify the path of the directory containing `./install.sh` as your IBM Tivoli Monitoring home directory. On certain platforms, this can cause the plug-in JAR files to overwrite themselves and become zero length files. The installation will fail as a result.

In addition, do not use the number sign (#) or exclamation mark (!) characters anywhere in your installation directory path.

Results

The Warehouse Proxy Agent or Summarization and Pruning Agent should now be upgraded and should not change the previous configuration. Manage Tivoli Monitoring Services should show the new version of the Warehouse Proxy Agent or Summarization and Pruning Agent.

The config files `hd.ini` as well as `sy.ini` should still contain the same information that it had before the upgrade, in addition to the new parameters introduced in the current release. The Warehouse Proxy Agent or Summarization and Pruning Agent can be restarted without reconfiguration, provided they were configured correctly before upgrade.

Note: The Warehouse Proxy Agent and Summarization and Pruning Agent are not started automatically after the upgrade has completed. Depending on the installed agents and the associated attribute groups enabled for historical collection, and whether limited database permissions are granted to the warehouse user, a database administrator might need to use the schema publication tool to generate a script with the required changes for the database. For more information, see Chapter 20, “Schema Publication Tool,” on page 525.

Upgrading the Warehouse Proxy Agent and Summarization and Pruning Agent using remote deploy

You can upgrade the Warehouse Proxy Agent or the Summarization and Pruning Agent by remotely deploying bundles containing the installation image and other prerequisites.

About this task

Complete the following steps to upgrade the Warehouse Proxy Agent and Summarization and Pruning Agent using remote deploy:

Procedure

1. For remote management of monitoring agents in general, you must install the appropriate OS monitoring agent. For instance, to remotely upgrade a Warehouse Proxy Agent or Summarization and Pruning Agent on Linux, you must install the monitoring agent for Linux OS on the same server. To remotely upgrade a Warehouse Proxy Agent or Summarization and Pruning Agent on Windows, you must install the monitoring agent for Windows on the same server.
2. Populate the DEPOT with the latest version of the Warehouse Proxy Agent or the Summarization and Pruning Agent. The DEPOT must be located where a Tivoli Enterprise Monitoring Server is installed. The CD image contains a directory called deploy where the Warehouse Proxy Agent bundle and the Summarization and Pruning Agent bundle are located. From this directory you can run the following commands:

Warehouse Proxy Agent command-line example:

```
tacmd addbundles -i . -t hd
```

Summarization and Pruning Agent command-line example:

```
tacmd addbundles -i . -t sy
```

Note: You can also use the install command `install.sh` to populate the depot with the latest version of the Warehouse Proxy Agent or the Summarization and Pruning Agent.

3. Connect to the Tivoli Enterprise Monitoring Server using the `tacmd` command:

```
tacmd login -s <TEMS hostname> -u <userid> -p <pw>
```

4. Run the update command to update to the latest version.

Warehouse Proxy Agent:

```
./tacmd updateAgent -n <system name> -t HD
```

Summarization and Pruning Agent:

```
./tacmd updateAgent -n <system name> -t SY
```

Results

The Warehouse Proxy Agent or the Summarization and Pruning Agent should now be updated but not restarted. A new version should appear in the topology workspace, the manage system status, manage monitoring services, and in the trace file.

You should now restart the Warehouse Proxy Agent or the Summarization and Pruning Agent, and check that the configuration has not been modified. The agents should start successfully.

Warehouse Proxy Agent silent install response file sample

```
INSTALL_ENCRYPTION_KEY=IBMTivoliMonitoringEncryptionKey
INSTALL_PRODUCT=hd
```

Warehouse Proxy Agent silent config response file sample

```
##### PRIMARY TEMS CONFIGURATION #####
# Will this agent connect to a Tivoli Enterprise Monitoring Server (TEMS)?
# This parameter is required.
# Valid values are: YES and NO
CMSCONNECT=YES

# What is the hostname of the TEMS to connect to?
# This parameter is NOT required. (default is the local system hostname)
```

```

# The TEMS must be the HUB TEMS.
HOSTNAME=localhost

# Will this agent connect to the TEMS through a firewall?
# This parameter is NOT required. (default is NO)
# Valid values are: YES and NO
# - If set to YES the NETWORKPROTOCOL must be ip.pipe
#FIREWALL=NO

# What network protocol is used when connecting to the TEMS?
# This parameter is required.
# Valid values are: ip, sna, ip.pipe, or ip.spice
NETWORKPROTOCOL=ip.pipe

# What is the first backup network protocol used for connecting to the TEMS?
# This parameter is NOT required. (default is none)
# Valid values are: ip, sna, ip.pipe, ip.spice, or none
#BK1NETWORKPROTOCOL=none

# What is the second backup network protocol used for connecting to the TEMS?
# This parameter is NOT required. (default is none)
# Valid values are: ip, sna, ip.pipe, ip.spice or none
#BK2NETWORKPROTOCOL=none

# If ip.pipe is one of the three protocols what is the IP pipe port number?
# This parameter is NOT required. (default is 1918)
IPPIPEPORTNUMBER=1918

# If ip.pipe is one of the three protocol what is the IP pipe partition name?
# This parameter is NOT required. (default is null)
#KDC_PARTITIONNAME=null

# If ip.pipe is one of the three protocols what is the KDC partition file?
# This parameter is NOT required. (default is null)
#KDC_PARTITIONFILE=null

# If ip.spice is one of the three protocols what is the IP pipe port number?
# This parameter is NOT required. (default is 3660)
#IPSPIPEPORTNUMBER=3660

# If ip is one of the three protocols what is the IP port number?
# This parameter is NOT required. (default is 1918)
# A port number and or one or more pools of port numbers can be given.
# The format for a pool is #-# with no embedded blanks.
#PORTNUMBER=1918

# If sna is one of the three protocols what is the SNA net name?
# This parameter is NOT required. (default is CANDLE)
#NETNAME=CANDLE

# If sna is one of the three protocols what is the SNA LU name?
# This parameter is NOT required. (default is LUNAME)
#LUNAME=LUNAME

# If sna is one of the three protocols what is the SNA log mode?
# This parameter is NOT required. (default is LOGMODE)
#LOGMODE=LOGMODE

##### SECONDARY TEMS CONFIGURATION #####

# Would you like to configure a connection for a secondary TEMS?
# This parameter is NOT required. (default is NO)
# Valid values are: YES and NO
#FTO=NO

# If configuring a connection for a secondary TEMS, what is the hostname of the secondary TEMS?
# This parameter is required if FTO=YES

```

```

#MIRROR=somehost.somewhere.com

# Will the agent connect to the secondary TEMS through a firewall?
# This parameter is NOT required. (default is NO)
# Valid values are: YES and NO
#FIREWALL2=NO

# What network protocol is used when connecting to the secondary TEMS?
# This parameter is required when FTO=YES and FIREWALL2 is NO
# Valid values are: ip, sna, or ip.pipe
#HSNETWORKPROTOCOL=ip.pipe

# What is the first backup network protocol used for connecting to the secondary TEMS?
# This parameter is NOT required. (default is none)
# Valid values are: ip, sna, ip.pipe, or none
#BK1HSNETWORKPROTOCOL=none

# What is the second backup network protocol used for connecting to the secondary TEMS?
# This parameter is NOT required. (default is none)
# Valid values are: ip, sna, ip.pipe, or none
#BK2HSNETWORKPROTOCOL=none

# If ip.pipe is one of the three secondary TEMS protocols what is the IP pipe port number?
# This parameter is NOT required. (default is 1918)
#HSIPPIPEPORTNUMBER=1918

# If ip is one of the three secondary TEMS protocols what is the IP port number?
# This parameter is NOT required. (default is 1918)
# A port number and or one or more pools of port numbers can be given.
# The format for a pool is #-# with no embedded blanks.
#HSPORTNUMBER=1918

# If sna is one of the three secondary TEMS protocols what is the SNA net name?
# This parameter is NOT required. (default is CANDLE)
#HSNETNAME=CANDLE

# If sna is one of the three secondary TEMS protocols what is the SNA LU name?
# This parameter is NOT required. (default is LUNAME)
#HSLUNAME=LUNAME

# If sna is one of the three secondary TEMS protocols what is the SNA log mode?
# This parameter is NOT required. (default is LOGMODE)
#HSLOGMODE=LOGMODE

##### OPTIONAL PRIMARY NETWORK NAME CONFIGURATION #####

# If the system is equipped with dual network host adapter cards you can designate
# another network name. What is the network name?
# This parameter is NOT required. (default is none)
#PRIMARYIP=none

#
#-----*
# WPA Specific Configuration Values
#-----*
# Allowed values are DB2, ORACLE, MSSQL
KHD_DBMS=DB2

#Comma separated list of fully qualified paths to JDBC JAR files
KHD_WAREHOUSE_JARS=/data/ana-db2/db2jcc.jar,/data/ana-db2/db2jcc_license_cu.jar

#-----*
# Specify the URL that corresponds to the warehouse type you selected
#
# Oracle URL
#KHD_ORACLE_JDBCURL=jdbc:oracle:thin:@<server>:<port>:<database>

```

```

#
# Microsoft SQL Seerver URL
#KHD_MSSQL_JDBCURL=jdbc:sqlserver://<server>:<port>;datasname=<database>;SelectMethod=cursor
#
# DB2 URL
KHD_DB2_JDBCURL=jdbc:db2://localhost:60000/ITMDW
#-----*

#-----*
# Specify the JDBC driver class that corresponds to the warehouse type you selected
#
# DB2 driver class
KHD_DB2_JBCDRIVER=com.ibm.db2.jcc.DB2Driver
#
# Oracle driver class
#KHD_ORACLE_JBCDRIVER=oracle.jdbc.driver.OracleDriver
#
# Microsoft SQL Server driver class
#KHD_MSSQL_JBCDRIVER=com.microsoft.sqlserver.jdbc.SQLServerDriver
#-----*

KHD_WAREHOUSE_USER=itmuser
KHD_WAREHOUSE_PASSWORD=itmtw08

# Should inserts be batched
# allowed values are true/false
KHD_BATCH_USE=true

# Enable Database Compression
# allowed values are true/false
KHD_DB_COMPRESSION=false

# List of TEMS served by this Warehouse Proxy
KHD_WAREHOUSE_TEMS_LIST=HUB_ITMTDWP12

# Enable Warehouse Compression for Z sources
# allowed values are true/false
KHD_SERVER_Z_COMPRESSION_ENABLE=false

# Enable Warehouse Compression for Distributed sources
# allowed values are true/false
KHD_SERVER_DIST_COMPRESSION_ENABLE=true

```

Summarization and Pruning Agent silent install response file sample

```

INSTALL_ENCRYPTION_KEY=IBMTivoliMonitoringEncryptionKey

INSTALL_PRODUCT=sy

```

Summarization and Pruning Agent silent config response file sample

```

# This is a sample silent configuration response file for the summarization
# and pruning agent

# To configure an agent using this silent response file:
# 1) copy this file to another location and modify the necessary parameters
# 2) run itmcmd config -A -p <silent_response_file> sy
#    - give an absolute path for the silent_response_file

# Syntax rules:
# - Comment lines begin with "#"
# - Blank lines are ignored
# - Parameter lines are PARAMETER=value (do not put space before the PARAMETER)
# - Space before or after an equal sign is OK

```

```

# - Parameter values should NOT contain the following characters $, =, or |

##### PRIMARY TEMS CONFIGURATION #####

# Will this agent connect to a Tivoli Enterprise Monitoring Server (TEMS)?
# This parameter is required.
# Valid values are: YES and NO
CMSCONNECT=YES

# What is the hostname of the TEMS to connect to?
# This parameter is NOT required. (default is the local system hostname)
HOSTNAME=localhost

# Will this agent connect to the TEMS through a firewall?
# This parameter is NOT required. (default is NO)
# Valid values are: YES and NO
# - If set to YES the NETWORKPROTOCOL must be ip.pipe
#FIREWALL=NO

# What network protocol is used when connecting to the TEMS?
# This parameter is required.
# Valid values are: ip, sna, ip.pipe, or ip.spipe
NETWORKPROTOCOL=ip.pipe

# What is the first backup network protocol used for connecting to the TEMS?
# This parameter is NOT required. (default is none)
# Valid values are: ip, sna, ip.pipe, ip.spipe, or none
#BK1NETWORKPROTOCOL=none

# What is the second backup network protocol used for connecting to the TEMS?
# This parameter is NOT required. (default is none)
# Valid values are: ip, sna, ip.pipe, ip.spipe or none
#BK2NETWORKPROTOCOL=none

# If ip.pipe is one of the three protocols what is the IP pipe port number?
# This parameter is NOT required. (default is 1918)
IPPIPEPORTNUMBER=1918

# If ip.pipe is one of the three protocol what is the IP pipe partition name?
# This parameter is NOT required. (default is null)
#KDC_PARTITIONNAME=null

# If ip.pipe is one of the three protocols what is the KDC partition file?
# This parameter is NOT required. (default is null)
#KDC_PARTITIONFILE=null

# If ip.spipe is one of the three protocols what is the IP pipe port number?
# This parameter is NOT required. (default is 3660)
#IPSPIPEPORTNUMBER=3660

# If ip is one of the three protocols what is the IP port number?
# This parameter is NOT required. (default is 1918)
# A port number and or one or more pools of port numbers can be given.
# The format for a pool is #-# with no embedded blanks.
#PORTNUMBER=1918

# If sna is one of the three protocols what is the SNA net name?
# This parameter is NOT required. (default is CANDLE)
#NETNAME=CANDLE

# If sna is one of the three protocols what is the SNA LU name?
# This parameter is NOT required. (default is LUNAME)
#LUNAME=LUNAME

# If sna is one of the three protocols what is the SNA log mode?
# This parameter is NOT required. (default is LOGMODE)
#LOGMODE=LOGMODE

```



```

##### SECONDARY TEMS CONFIGURATION #####

# Would you like to configure a connection for a secondary TEMS?
# This parameter is NOT required. (default is NO)
# Valid values are: YES and NO
#FTO=NO

# If configuring a connection for a secondary TEMS, what is the hostname of the secondary TEMS?
# This parameter is required if FTO=YES
#MIRROR=somehost.somewhere.com

# Will the agent connect to the secondary TEMS through a firewall?
# This parameter is NOT required. (default is NO)
# Valid values are: YES and NO
#FIREWALL2=NO

# What network protocol is used when connecting to the secondary TEMS?
# This parameter is required when FTO=YES and FIREWALL2 is NO
# Valid values are: ip, sna, or ip.pipe
#HSNETWORKPROTOCOL=ip.pipe

# What is the first backup network protocol used for connecting to the secondary TEMS?
# This parameter is NOT required. (default is none)
# Valid values are: ip, sna, ip.pipe, or none
#BK1HSNETWORKPROTOCOL=none

# What is the second backup network protocol used for connecting to the secondary TEMS?
# This parameter is NOT required. (default is none)
# Valid values are: ip, sna, ip.pipe, or none
#BK2HSNETWORKPROTOCOL=none

# If ip.pipe is one of the three secondary TEMS protocols what is the IP pipe port number?
# This parameter is NOT required. (default is 1918)
#HSIPPIPEPORTNUMBER=1918

# If ip is one of the three secondary TEMS protocols what is the IP port number?
# This parameter is NOT required. (default is 1918)
# A port number and or one or more pools of port numbers can be given.
# The format for a pool is #-# with no embedded blanks.
#HSPORTNUMBER=1918

# If sna is one of the three secondary TEMS protocols what is the SNA net name?
# This parameter is NOT required. (default is CANDLE)
#HSNETNAME=CANDLE

# If sna is one of the three secondary TEMS protocols what is the SNA LU name?
# This parameter is NOT required. (default is LUNAME)
#HSLUNAME=LUNAME

# If sna is one of the three secondary TEMS protocols what is the SNA log mode?
# This parameter is NOT required. (default is LOGMODE)
#HSLOGMODE=LOGMODE

##### OPTIONAL PRIMARY NETWORK NAME CONFIGURATION #####

# If the system is equipped with dual network host adapter cards you can designate
# another network name. What is the network name?
# This parameter is NOT required. (default is none)
#PRIMARYIP=none

#
#-----*
# SNP Specific Configuration Values
#-----*
# Allowed values are DB2, ORACLE, MSSQL

```

```

KSY_WAREHOUSE_TYPE=DB2

#Comma separated list of fully qualified paths to JDBC JAR files
KSY_WAREHOUSE_JARS=

#-----*
# Specify the URL that corresponds to the warehouse type you selected
#
# Oracle URL
#KSY_ORACLE_JDBCURL=jdbc:oracle:thin:@<server>:<port>:<database>
#
# Microsoft SQL Seerver URL
#KSY_MSSQL_JDBCURL=jdbc:sqlserver://<server>:<port>;databasename=<database>;SelectMethod=cursor
#
# DB2 URL
KSY_DB2_JDBCURL=jdbc:db2://<server>:<port>/<database>
#-----*

#-----*
# Specify the JDBC driver class that corresponds to the warehouse type you selected
#
# DB2 driver class
KSY_DB2_JDBCDRIVER=com.ibm.db2.jcc.DB2Driver
#
# Oracle driver class
#KSY_ORACLE_JDBCDRIVER=oracle.jdbc.driver.OracleDriver
#
# Microsoft SQL Server driver class
#KSY_MSSQL_JDBCDRIVER=com.microsoft.sqlserver.jdbc.SQLServerDriver
#-----*

KSY_WAREHOUSE_USER=
KSY_WAREHOUSE_PASSWORD=

#-----*
# Warehouse Database Compression                                *
# If set to Y, database compression will be enabled.            *
#-----*
KSY_DB_COMPRESSION=N
#
#-----*
# Timezone indicator                                            *
# AGENT means use the time zone offset of the agent            *
# WAREHOUSE means use the time zone offset of the warehouse    *
#-----*
KSY_TIMEZONE_IND=AGENT
#
#-----*
# Start of the week day                                        *
# 0 = Sunday 1 = Monday                                        *
#-----*
KSY_START_OF_WEEK_DAY=0

#
#-----*
# Shift periods                                                *
# Only two shifts are allowed.                                  *
# If shifts are enabled,                                        *
# each hour (0-23) must be specified once                      *
#-----*
KSY_SHIFTS_ENABLED=N
KSY_SHIFT1_HOURS=0,1,2,3,4,5,6,7,8,18,19,20,21,22,23
KSY_SHIFT2_HOURS=9,10,11,12,13,14,15,16,17
#
#-----*
# Vacation periods                                             *

```

```

# Vacation days are comma delimited and use the YYYYMMDD format *
#-----*
KSY_VACATIONS_ENABLED=N
KSY_WEEKENDS_AS_VACATIONS=N
KSY_VACATION_DAYS=
#
#-----*
# Maximum rows per database transaction *
#-----*
KSY_MAX_ROWS_PER_TRANSACTION=1000
#
#-----*
# Schedule *
# Use KSY_FIXED_SCHEDULE to indicate you would like the S&P agent *
# to run at a set time every day or to run periodically *
# throughout the day. *
# KSY_FIXED_SCHEDULE=Y means the S&P agent will use the *
# KSY_EVERY_N_DAYS, KSY_HOUR_TO_RUN, KSY_HOUR_AM_PM and *
# KSY_MINUTE_TO_RUN variables to determine at which fixed time *
# to run. *
# KSY_FIXED_SCHEDULE=N means the S&P agent will use the *
# KSY_EVERY_N_MINS and KSY_BLACKOUT variables to determine *
# how often and when to run. *
#
# KSY_BLACKOUT can be used with KSY_FIXED_SCHEDULE=N to list *
# ranges of time where the S&P agent should not run. The list is *
# in the format HH:MM-HH:MM with multiple values separated by a *
# comma. The start time must be lower than the end time and the *
# hour values must be between 0 and 23. To have S&P agent run *
# every hour except between 12 AM and 2 AM, use the following: *
# KSY_BLACKOUT=00:00-01:59 *
#-----*
KSY_FIXED_SCHEDULE=Y
KSY_EVERY_N_DAYS=1
KSY_HOUR_TO_RUN=2
KSY_HOUR_AM_PM=AM
KSY_MINUTE_TO_RUN=0
KSY_EVERY_N_MINS=60
KSY_BLACKOUT=

#-----*
# KSY_SUMMARIZATION_UNITS and KSY_NODE_ERROR_UNITS determine the *
# number of rows that will be shown in the Summarization and *
# Pruning overall run table and Errors table respectively. There *
# is a minimum value of 10 for each *
#-----*
KSY_SUMMARIZATION_UNITS=10
KSY_NODE_ERROR_UNITS=10
#
#-----*
# KSY_BATCH_MODE controls the batching method used by the *
# Summarization and Pruning agent. *
# A value of 0 means that data should only be batched for the same *
# system while a value of 1 means that data from multiple systems *
# can be batched together which can lead to higher performance at *
# potentially bigger transaction sizes. Default is 0 (single). *
#-----*
KSY_BATCH_MODE=0
#
#-----*
# WAREHOUSELOG and WAREHOUSEAGGREGLOG pruning settings *
# The entries are in the format <number>.<unit> where <number> *
# is the number of units to keep and <unit> is one of: *
# day, month or year. For example, keep data for 14 days, set the *
# value to 14.day *
#-----*
KSY_WAREHOUSELOG_PRUNE=

```

```

KSY_WAREHOUSEAGGREGLOG_PRUNE=
#
#-----*
# CNP Server connection *
# default host = localhost *
# default port = 1920 *
#-----*
KSY_CNP_SERVER_HOST=localhost
KSY_CNP_SERVER_PORT=1920
#
#-----*
# Age Units *
# Defines the minimum age of data before aggregation is done *
# Only applies to the lowest aggregation unit for a product *
# Only HOUR and DAY are supported *
#-----*
KSY_HOUR_AGE_UNITS=1
KSY_DAY_AGE_UNITS=0

#-----*
# Maximum number of simultaneous worker threads *
# Default is 2 *
# Recommended value: number of processors on your server times 2 *
#-----*
KSY_MAX_WORKER_THREADS=2

#-----*
# KSY_CACHE_MINS records the number of minutes to cache the *
# database connectivity for MOSWOS reporting purposes. The *
# minimum value is 5 minutes. *
#-----*
KSY_CACHE_MINS=10

```

Configuring the Warehouse Proxy Agent using the Tivoli Enterprise Portal

This function allows you to remotely configure the Warehouse Proxy Agent using the Tivoli Enterprise Portal. You can use this function anytime the Warehouse Proxy Agent configuration needs to be changed.

For remote management of monitoring agents in general, you must install the appropriate OS monitoring agent. For instance, to remotely manage a Warehouse Proxy Agent on Linux systems, you must install the Monitoring Agent for Linux OS on the same server. To remotely manage a Warehouse Proxy Agent on Windows systems, you must install the Monitoring Agent for Windows on the same server.

Also, to remotely configure the agent, the Warehouse Proxy Agent itself must be added to the remote depot. The CD image contains a deploy directory where the Warehouse Proxy Agent bundle is located. From this directory you can run the following command:

```
tacmd addbundles -i . -t hd
```

Note:

1. The **Test Connection** button is not available when configuring through the Tivoli Enterprise Portal.
2. The Warehouse Proxy Agent for the target platform being configured must have already been added to the Tivoli Enterprise Monitoring Server depot for the configuration to succeed.
3. Browsing for JDBC JAR files is not supported since they are on a remote system. You must enter the comma separated, fully qualified list of JDBC JAR files. The same as you do in the CLI.

The Deployment Status Detail workspace should show a SUCCESSFUL Transaction for the Configure command against the Warehouse Proxy Agent. The configuration file on the remote system where

Warehouse Proxy Agent is installed should contain the corresponding changes executed on the configuration panel. Configuration parameters for the Warehouse Proxy Agent are set in the following files according to operating system:

Windows

ITM_HOME\TMAITM6\khdenv

For example: C:\IBM\ITM\TMAITM6\khdenv

Linux

UNIX

ITM_HOME/config/hd.ini

For example: /opt/IBM/ITM/config/hd.ini

Afterwards, the Warehouse Proxy Agent operates as usual when configuration is successful. Connectivity of the Warehouse Proxy Agent can be checked on the Warehouse Proxy Agent Status workspace.

To enable this function, complete the following steps:

1. Open the Tivoli Enterprise Portal.
2. Right-click on the Warehouse Proxy Agent in the navigator tree.
3. Choose **Configure**. The configuration panel for the Warehouse Proxy Agent is displayed.

The log file can be found in the %CANDLEHOME%/cnp/logs directory.

Configuring the Summarization and Pruning Agent using the Tivoli Enterprise Portal

This function allows you to remotely configure the Summarization and Pruning Agent using the Tivoli Enterprise Portal. You can use this function anytime the Summarization and Pruning Agent configuration needs to be changed.

Before you begin

For remote management of monitoring agents in general, you must install the appropriate OS monitoring agent. For instance, to remotely manage a Summarization and Pruning Agent on Linux systems, you must install the Monitoring Agent for Linux OS on the same server. To remotely manage a Summarization and Pruning Agent on Windows systems, you must install the Monitoring Agent for Windows on the same server.

Also, to remotely configure the agent, the Summarization and Pruning Agent itself must be added to the remote depot. The CD image contains a deploy directory where the Summarization and Pruning Agent bundle is located. From this directory you can run the following command:

```
tacmd addbundles -i . -t sy
```

About this task

The Deployment Status Detail workspace should show a SUCCESSFUL Transaction for the Configure command against the Summarization and Pruning Agent. The configuration file (sy.ini or KSYENV) on the remote system where the Summarization and Pruning Agent is installed should contain the corresponding changes executed on the configuration panel. Afterwards, the Summarization and Pruning Agent operates as usual when configuration is successful. Connectivity of the Summarization and Pruning Agent can be checked on the Summarization and Pruning Status workspace.

To enable this function, complete the following steps:

Procedure

1. Open the Tivoli Enterprise Portal.
2. Select the Summarization and Pruning Agent that needs to be remotely configured.
3. Right-click on the Summarization and Pruning Agent in the navigator tree.
4. Choose **Configure**. The configure panel for the Summarization and Pruning Agent is displayed.

Results

The Tivoli Enterprise Portal log file can be found in the %CANDLEHOME%\cnp\logs directory for Windows systems, or the installer/CANDLEHOME/logs directory for UNIX systems. Remote deploy tracing can also be set: ERROR(UNIT: KDY ALL) on the OS agent, ERROR(UNIT: KDY ALL) on the monitoring server.

Note:

1. The **Test Connection** button is not available when configuring through the Tivoli Enterprise Portal.
2. The Summarization and Pruning Agent for the target platform being configured must have already been added to the Tivoli Enterprise Monitoring Server depot for the configuration to succeed.
3. Browsing for JDBC JAR files is not supported since they are on a remote system. You must enter the comma separated, fully qualified list of JDBC JAR files. The same as you do in the CLI.

Remotely starting and stopping the Warehouse Proxy Agent using the Tivoli Enterprise Portal

This function allows you to remotely manage the Warehouse Proxy Agent using the Tivoli Enterprise Portal. You can use this function anytime the Warehouse Proxy Agent needs to be started or stopped.

Before you begin

For remote management of monitoring agents in general, you must install the appropriate OS monitoring agent. For instance, to remotely manage a Warehouse Proxy Agent on Linux systems, you must install the Monitoring Agent for Linux OS on the same server. To remotely manage a Warehouse Proxy Agent on Windows systems, you must install the Monitoring Agent for Windows on the same server

About this task

The Deployment Status Detail workspace should show a SUCCESSFUL Transaction for the Configure command against the Warehouse Proxy Agent. The Warehouse Proxy Agent in the navigator tree should be grayed out if the agent is stopped, and visible if it is started. Afterwards, the Warehouse Proxy Agent operates as usual when configuration is successful. Connectivity of the Warehouse Proxy Agent can be checked on the Warehouse Proxy Agent Status workspace.

To enable this function, complete the following steps:

Procedure

1. Open the Tivoli Enterprise Portal.
2. Select the Warehouse Proxy Agent that needs to be remotely managed.
3. Right-click on the Warehouse Proxy Agent in the navigator tree.
4. Choose **Start** or **Stop**. The configure panel for the Warehouse Proxy Agent is displayed.

Results

The Tivoli Enterprise Portal log file can be found in the %CANDLEHOME%\cnp\logs directory for Windows systems, or the installer/CANDLEHOME/logs directory for UNIX systems. Remote deploy tracing can also be set: ERROR(UNIT: KDY ALL) on the OS agent, ERROR(UNIT: KDY ALL) on the monitoring server.

Remotely starting and stopping the Summarization and Pruning Agent using the Tivoli Enterprise Portal

This function allows you to remotely manage the Summarization and Pruning Agent using the Tivoli Enterprise Portal. You can use this function anytime the Summarization and Pruning Agent needs to be started or stopped.

Before you begin

For remote management of monitoring agents in general, you must install the appropriate OS monitoring agent. For instance, to remotely manage a Summarization and Pruning Agent on Linux systems, you must install the Monitoring Agent for Linux OS on the same server. To remotely manage a Summarization and Pruning Agent on Windows systems, you must install the Monitoring Agent for Windows on the same server.

About this task

The Deployment Status Detail workspace should show a SUCCESSFUL Transaction for the Configure command against the Summarization and Pruning Agent. The Summarization and Pruning Agent in the navigator tree should be grayed out if the agent is stopped, and visible if it is started. Afterwards, the Summarization and Pruning Agent operates as usual when configuration is successful. Connectivity of the Summarization and Pruning Agent can be checked on the Summarization and Pruning Agent Status workspace.

To enable this function, complete the following steps:

Procedure

1. Open the Tivoli Enterprise Portal.
2. Select the Summarization and Pruning Agent that needs to be remotely managed.
3. Right-click on the Summarization and Pruning Agent in the navigator tree.
4. Choose **Start** or **Stop**. The configure panel for the Summarization and Pruning Agent is displayed.

Results

The Tivoli Enterprise Portal log file can be found in the %CANDLEHOME%\cnp\logs directory for Windows systems, or the installer/CANDLEHOME/logs directory for UNIX systems. Remote deploy tracing can also be set: ERROR(UNIT: KDY ALL) on the OS agent, ERROR(UNIT: KDY ALL) on the monitoring server.

Remotely deploying the Warehouse Proxy Agent

This function allows you to remotely deploy the Warehouse Proxy Agent. You can use this function anytime the Warehouse Proxy Agent needs to be installed remotely.

For remote management of monitoring agents in general, you must install the appropriate OS monitoring agent. For instance, to remotely manage a Warehouse Proxy Agent on Linux systems, you must install the Monitoring Agent for Linux OS on the same server. To remotely manage a Warehouse Proxy Agent on Windows systems, you must install the Monitoring Agent for Windows on the same server.

Also, to remotely deploy the Warehouse Proxy Agent, you must populate the remote depot with the necessary bundles (the Warehouse Proxy Agent bundle and its prerequisites) for the operating systems where the Warehouse Proxy Agent needs to be deployed. This can be done using the product installer or the **tacmd addBundles** command.

During this process of adding bundles, IBM Tivoli Monitoring reads the descriptor file (dsc) of the specified bundle you are adding, and identifies (among all of the files in the installation media) which files need to be copied on the agent's depot in your monitoring server. The files loaded are copied onto the path specified in the DEPOTHOME environment variable, defined in the KBBENV environment file. The **tacmd viewDepot** command allows you to display the types of agents in the deployment depot. The CD image contains a deploy directory where the Warehouse Proxy Agent bundle is located. When the depot is populated, you must deploy an OS agent on the remote computer, using the **tacmd createNode** command.

After the node has been created and OS agent is running, you can install other non-OS agents. This installation can be done by using the Tivoli Enterprise Portal or by using the **tacmd createNode** command.

Run the following command to complete enabling this function:

```
tacmd addSystem -t hd -n Os_node_name
```

Properties supported:

KHD_DB_TYPE.KHD_DBMS
Allowed values are DB2, ORACLE, MSSQL
KHD_PARMS.KHD_BATCH_USE
Controls if batching is enabled. Allowed values are true, false. Default is true.
KHD_PARMS.KHD_DB_COMPRESSION
Controls if database compression should be used when creating tables and indexes.
Allowed values are true, false. Default is false.
KHD_PARMS.KHD_WAREHOUSE_TEMS_LIST
Comma-separated list of TEMS that this Warehouse Proxy is responsible for.
KHD_PARMS.KHD_WAREHOUSE_USER
The user ID used to connect to TDW
KHD_PARMS.KHD_WAREHOUSE_PASSWORD
The password used to connect to TDW.
KHD_PARMS.KHD_SERVER_Z_COMPRESSION_ENABLE
Allowed values are true, false. Default is false.
KHD_PARMS.KHD_SERVER_DIST_COMPRESSION_ENABLE
Allowed values are true, false. Default is true.

Windows Only:

KHD_PARMS.KHD_ODBC_DSN
The ODBC data source name to be used.

Unix/Linux only:

KHD_PARMS.KHD_DB2_JDBCURL
The JDBC URL to connect to DB2. Default is jdbc:db2://localhost:50000/WAREHOUS
KHD_PARMS.KHD_DB2_JDBCDRIVER
The DB2 JDBC driver. Default is com.ibm.db2.jcc.DB2Driver
KHD_PARMS.KHD_ORACLE_JDBCURL
The JDBC URL to connect to URL. Default is jdbc:oracle:thin:@<server>:<port>:<database>
KHD_PARMS.KHD_ORACLE_JDBCDRIVER
The Oracle JDBC drive. Default is oracle.jdbc.driver.OracleDriver
KHD_PARMS.KHD_MSSQL_JDBCURL
The JDBC URL to connect to Microsoft SQL Server. Default is
jdbc:sqlserver://<server>:<port>;datasource=<database>;
SelectMethod=cursor
KHD_PARMS.KHD_MSSQL_JDBCDRIVER
The Microsoft SQL Server JDBC driver. The default is
com.microsoft.sqlserver.jdbc.SQLServerDriver
KHD_PARMS.KHD_WAREHOUSE_JARS
The comma separated list of JDBC JARs that are needed to connect to the database.

A Warehouse Proxy Agent should be installed and configured on the remote system once the transaction has completed successfully. The Warehouse Proxy Agent operates as usual when it is started. Connectivity of the Warehouse Proxy Agent can be checked on the Warehouse Proxy Agent Status workspace.

The `configureSystem` command allows CLI configuration of a remote agent. If `KHD_PARMS.KHD_WAREHOUSE_TEMS_LIST` is not provided, the value will be set to an empty value.

It is the OS Agent that is used by remote deploy to do the work. The OS Agent log on the endpoint shows all the remote deploy activity (kdy component) processing. If you enable `ERROR(UNIT:KDY ALL)` on the endpoint's OS Agent tracing, you will see the remote deploy processing. To redeploy the same agent, run the **`tacmd updateagent`** command using the `-v` option.

Note: Remote deploy will start the agent after installation which might cause situations to fire due to lack of default configuration. The configuration supplied during the remote deploy is applied after the install and start is done. The agent will be restarted after the configuration is complete.

Remotely deploying the Summarization and Pruning Agent

This function allows you to remotely deploy the Summarization and Pruning Agent. You can use this function anytime the Summarization and Pruning Agent needs to be installed remotely.

For remote management of monitoring agents in general, you must install the appropriate OS monitoring agent. For instance, to remotely manage a Summarization and Pruning Agent on Linux systems, you must install the Monitoring Agent for Linux OS on the same server. To remotely manage a Summarization and Pruning Agent on Windows systems, you must install the Monitoring Agent for Windows on the same server.

Also, to remotely deploy the Summarization and Pruning Agent, you must populate the remote depot with the necessary bundles (the Summarization and Pruning Agent bundle and its prerequisites) for the operating systems where the Summarization and Pruning Agent needs to be deployed. This can be done using the product installer or the **`tacmd addBundles`** command.

During this process of adding bundles, IBM Tivoli Monitoring reads the descriptor file (dsc) of the specified bundle you are adding, and identifies (among all of the files in the installation media) which files need to be copied on the agent's depot in your monitoring server. The files loaded are copied onto the path specified in the `DEPOTHOME` environment variable, defined in the `KBBENV` environment file. The **`tacmd viewDepot`** command allows you to display the types of agents in the deployment depot. The CD image contains a deploy directory where the Summarization and Pruning Agent bundle is located. When the depot is populated, you must deploy an OS agent on the remote computer, using the **`tacmd createNode`** command.

After the node has been created and OS agent is running, you can install other non-OS agents. This installation can be done by using the Tivoli Enterprise Portal or by using the **`tacmd createNode`** command.

Run the following command to complete enabling this function:

```
tacmd addSystem -t sy -n Os_node_name
```

Properties supported:

`DBTYPE.KSY_WAREHOUSE_TYPE`
The database type used for TDW. Allowed values are DB2, ORACLE, MSSQL

`SOURCES.KSY_WAREHOUSE_JARS`
The comma-separated list of JDBC JAR files needed to connect to the database.

`SOURCES.KSY_DB2_JDBCURL`
The JDBC URL used to connect to DB2 database. Default is `jdbc:db2://localhost:50000/WAREHOUS`

`SOURCES.KSY_DB2_JBCDRIVER`

The DB2 JDBC driver. Default is com.ibm.db2.jcc.DB2Driver
 SOURCES.KSY_ORACLE_JDBCURL
 The JDBC URL used to connect to Oracle. Default is jdbc:oracle:thin:@<server>:<port>:<database>
 SOURCES.KSY_ORACLE_JBCDRIVER
 The Oracle JDBC driver. Default is oracle.jdbc.driver.OracleDriver
 SOURCES.KSY_MSSQL_JDBCURL
 The JDBC URL used to connect to Microsoft SQL Server. Default is
 jdbc:sqlserver://<server>:<port>;datasname=<database>;SelectMethod=cursor
 SOURCES.KSY_MSSQL_JBCDRIVER
 The Microsoft SQL Server JDBC Driver. Default is com.microsoft.sqlserver.jdbc.SQLServerDriver
 SOURCES.KSY_WAREHOUSE_USER
 The user ID used to connect to the TDW database.
 SOURCES.KSY_WAREHOUSE_PASSWORD
 The password used to connect to the TDW database.
 SOURCES.KSY_CNP_SERVER_HOST
 The TEPS hostname. Default is localhost
 SOURCES.KSY_CNP_SERVER_PORT
 The TEPS port number. Default is 1920
 SCHEDULING.KSY_FIXED_SCHEDULE
 The scheduling to be used by S&P. Y for fixed schedule or N for flexible schedule.
 Default is Y
 SCHEDULING.KSY_EVERY_N_MINS
 The flexible scheduling interval in minutes. Default is 15
 SCHEDULING.KSY_EVERY_N_DAYS
 The fixed scheduling interval in days. Default is 1
 SCHEDULING.KSY_HOUR_TO_RUN
 The fixed scheduling time to run. Default is 2
 SCHEDULING.KSY_HOUR_AM_PM
 The fixed scheduling AM or PM of time to run. Allowed values: AM,
 PM Default is AM
 SCHEDULING.KSY_MINUTE_TO_RUN
 The fixed scheduling minute to run. Default is 0
 SCHEDULING.KSY_BLACKOUT
 Comma-separated list of blackout periods in the format HH:MM-HH:MM using
 24 hour format. The flexible scheduling will not start during the listed blackout periods.
 Note that the end time should be greater than the start time of the blackout period.
 WORK.KSY_START_OF_WEEK_DAY
 The start of the week day, used when computing weekly aggregates.
 Allowed values are 0 (Sunday), 1 (Monday). Default is 0 (Sunday)
 WORK.KSY_SHIFTS_ENABLED
 Controls whether shifts are used or not. Allowed values are Y or N Default is N
 WORK.KSY_SHIFT1_HOURS
 Comma-separated list of off peak shift hours. Default is
 0,1,2,3,4,5,6,7,8,18,19,20,21,22,23
 WORK.KSY_SHIFT2_HOURS
 Comma-separated list of peak shift hours. Default is 9,10,11,12,13,14,15,16,17
 WORK.KSY_VACATIONS_ENABLED
 Controls whether vacation days are used or not. Allowed values are Y or N Default is N
 WORK.KSY_WEEKENDS_AS_VACATIONS
 Controls whether weekends are treated as vacation days. Allowed values are Y or N.
 Default is N
 WORK.KSY_VACATION_DAYS
 Comma separated list of vacation days in the format YYYYMMDD where YYYY is the year,
 MM is the month (1-12) and DD is the day (1-31).
 LOG.KSY_WAREHOUSELOG_PRUNE
 Controls the pruning of the WAREHOUSELOG table. The format is nn.uuu where nn is
 the number of units to retain and uuu is one of day, month, or year.
 Example: 5.day to retain data for 5 days.
 LOG.KSY_WAREHOUSEAGGREGLOG_PRUNE
 Controls the pruning of the WAREHOUSEAGGREGLOG table. The format is nn.uuu where nn is
 the number of units to retain and uuu is one of day, month, or year.
 Example: 5.day to retain data for 5 days.
 ADDITIONAL.KSY_MAX_WORKER_THREADS
 The number of threads that will be used for summarization and pruning. Default is 2
 ADDITIONAL.KSY_MAX_ROWS_PER_TRANSACTION
 The number of rows per database transaction. Default is 1000
 ADDITIONAL.KSY_TIMEZONE_IND

Controls which time zone offset should be used. Allowed values are AGENT or WAREHOUSE Default is AGENT

ADDITIONAL.KSY_HOUR_AGE_UNITS
Controls the age in hours of data to be summarized. Default is 1

ADDITIONAL.KSY_DAY_AGE_UNITS
Controls the age in days of data to be summarized. Default is 0

ADDITIONAL.KSY_NODE_ERROR_UNITS
The number of node errors to keep for the self-monitoring workspace. Default is 10

ADDITIONAL.KSY_SUMMARIZATION_UNITS
The number of summarization runs to keep for the self monitoring workspace. Default is 10

ADDITIONAL.KSY_CACHE_MINS
The time in minutes that connectivity information is cached. Default is 15

ADDITIONAL.KSY_BATCH_MODE
Controls whether batching of data should be used to improve performance. Allowed values are 0 (single system) or 1 (multiple systems). Default is 0

ADDITIONAL.KSY_DB_COMPRESSION
Controls whether tables and indexes should be created with database compression enabled. Allowed values are N or Y. Default is N.

A Summarization and Pruning Agent should be installed and configured on the remote system once the transaction has completed successfully. The Summarization and Pruning Agent operates as usual when it is started. Connectivity of the Summarization and Pruning Agent can be checked on the Summarization and Pruning Agent Status workspace.

The configureSystem command allows CLI configuration of a remote agent. If SCHEDULING.KSY_BLACKOUT, WORK.KSY_VACATION_DAYS, LOG.KSY_WAREHOUSELOG_PRUNE or LOG.KSY_WAREHOUSEAGGREGLOG_PRUNE are not provided, their value will be set to an empty value.

It is the OS Agent that is used by remote deploy to do the work. The OS Agent log on the endpoint shows all the remote deploy activity (kdy component) processing. If you enable ERROR(UNIT:KDY ALL) on the endpoint's OS Agent tracing, you will see the remote deploy processing. To redeploy the same agent, run the **tacmd updateagent** command using the **-v** option.

Note: Remote deploy will start the agent after installation which may cause situations to fire due to lack of default configuration. The configuration supplied during the remote deploy is applied after the install and start is done. The agent will be restarted after the configuration is complete.

WAREHOUSELOG and WAREHOUSEAGGREGLOG tables

The WAREHOUSELOG table lets you know how many exports succeed and how many failed because of an ODBC error or a timeout value issue.

The WAREHOUSELOG table has the following columns. All times are the local time on the machine where the Warehouse Proxy instance is running.

ORIGINNODE

The name of the computer that made the request. This name is the node name for the agent. For example, Primary::box1:NT.

OBJECT

The attribute group that submitted the request. For example, NT_System.

STARTQUEUE

The time when the request was inserted in the work queue. For example, 10508201154000000.

ENDQUEUE

The time when the request exited the work queue. For example, 10508201155000000.

STARTEXPORT

The amount of time that elapsed before the first row of the sample request was retrieved. For example, 105082011562000000.

EXPORTTIME

The amount of time after the export request transaction was committed. For example, 10508201157000000.

ROWSINSERTED

The number of row inserted in the database for the request. For example, 1000.

ROWSRECEIVED

The number of rows retrieved from the RPC source. For example, 1000.

ROWSSKIPPED

This column is not used.

STARTTIME

The start time of the collection for that sample. For example, 10508150920000000.

ENDTIME

The end time of the collection for that sample. For example, 10508150920000000.

ERRORMSG

An error message when no rows are inserted in the database. The error message can indicate an ODBC error or a TIMEOUT error. For example:

Sample data rejected for timeout reason at stage COMMIT EXPORT

WPSYSNAME

The name of the Warehouse Proxy Agent that inserted the rows into the database.

The WAREHOUSEAGGREGLOG table logs the progress of the Summarization and Pruning Agent as it is processing data. Each time the Summarization and Pruning Agent executes, it adds an entry for each attribute group (OBJECT column) and origin node (ORIGINNODE column) that was processed. The WAREHOUSEAGGREGLOG table contains the following columns:

ORIGINNODE

The name of the computer that is being summarized. This name is the node name for the agent. For example, Primary::box1:NT. -

OBJECT

The attribute group that was processed.

LOGTMZDIFF

The time zone difference for the Summarization and Pruning Agent.

MINWRITETIME

The minimum WRITETIME value that was read from the sample data for the specified ORIGINNODE and OBJECT.

MAXWRITETIME

The maximum WRITETIME value that was read from the sample data for the specified ORIGINNODE and OBJECT

STARTTIME

The time that the Summarization and Pruning Agent processing began for the specified ORIGINNODE and OBJECT.

ENDTIME

The time that the Summarization and Pruning Agent processing ended for the specified ORIGINNODE and OBJECT.

ROWSREAD

The number of sample data rows read for the specified ORIGINNODE and OBJECT in the time interval MINWRITETIME and MAXWRITETIME.

Chapter 28. Using Tivoli Common Reporting and monitoring agent reports

The Tivoli Common Reporting tool (TCR) is a reporting feature available to users of Tivoli products and is a component of Jazz for Service Management. Use Tivoli Common Reporting to gather, analyze, and report important trends in your managed environment in a consistent and integrated manner. A set of predefined reports is provided for the Tivoli Monitoring OS Agents and other products for monitoring individual, multiple, and enterprise resources.

To start using reports complete the following steps:

1. Install Tivoli Common Reporting and its prerequisite software if it is not already in your environment. See “Required software and memory requirements for a reporting environment that uses Tivoli Common Reporting” on page 93 for more details.
2. Install and configure the Tivoli Data Warehouse and warehouse agents: Warehouse Proxy Agent and Summarization and Pruning Agent. For information on setting up Tivoli Data Warehouse and the warehouse agents, see Chapter 19, “Tivoli Data Warehouse solutions,” on page 509.
3. Configure historical collection. See *Historical collection configuration* in the *IBM Tivoli Monitoring Tivoli Enterprise Portal User's Guide*.
4. Prepare the Tivoli Data Warehouse database for reporting data. See *Maintaining the dimension tables* in the *IBM Tivoli Monitoring Administrator's Guide*.
5. Install the monitoring agent reports. See *Importing and running IBM Cognos® reports* in the *IBM Tivoli Monitoring Administrator's Guide*.
6. Review the reports. See the monitoring agent user's guide for report content specifics.

Documentation for the Tivoli Common Reporting tool is located at http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/index.jsp?topic=%2Fcom.ibm.psc.doc_1.1.0%2Finstall%2Ftcr_t_install.html.

Chapter 29. Integrating event management systems

If you are using either Netcool/OMNIBus or Tivoli Enterprise Console in addition to IBM Tivoli Monitoring to manage events in your enterprise, you can integrate and manage events from a single console. You can integrate event management by forwarding events reported by Tivoli Enterprise Monitoring Agents to either event system for correlation and management - changes in event status made on the event system are reflected back to the hub monitoring server that forwarded them. Or you can enable events reported by a Tivoli Enterprise Monitoring Agent to be passed directly to OMNIBus for processing, thereby bypassing the monitoring server entirely.

To enable forwarding of situation events, the Netcool/OMNIBus (the event server) or Tivoli Enterprise Console server must be configured to receive the events, a situation update forwarding process must be installed on the event server, situation forwarding must be enabled on either the hub monitoring server or the monitoring agent, and a default event integration facility (EIF) destination must be defined.

Using the Tivoli Enterprise Portal, you can define monitoring specifications, called *situations*, to detect the occurrence of specific conditions on managed systems. When conditions that match the specifications are detected by monitoring agents, events are reported to the monitoring servers. In environments where Netcool/OMNIBus or Tivoli Enterprise Console are also being used for event management, hub Tivoli Enterprise Monitoring Servers or the agents themselves can be configured to forward these situation events to either of these event servers for further correlation and management. Changes in the status of events made on the event server are reported back to the forwarding monitoring server so that events are synchronized on the two event management systems.

The two chapters that follow this section provide instructions for implementing event integration by forwarding situation events to Netcool/OMNIBus and Tivoli Enterprise Console.

In addition you can view events within the Tivoli Enterprise Portal. The *Common Event Console* is a Tivoli Enterprise Portal view that provides an integrated display of events from multiple event systems. The Common Event Console presents normalized events from the supported event systems in a single table. By default, the view displays events reported to a hub Tivoli Enterprise Monitoring Server by Tivoli Enterprise Monitoring Agents (*situation events*); it can also be configured to present events from Netcool/OMNIBus and Tivoli Enterprise Console event systems.

A common event connector (frequently called simply a connector) enables the integrated display of events from an event system in the Common Event Console. The connector retrieves event data from a supported event system and sends user-initiated actions to be run in that event system. To have the events from a specific event system displayed in the Common Event Console, you must configure a connector for that event system. Because the connector for Tivoli Enterprise Monitoring Agents is configured when you install the portal, the Common Event Console includes all situation events by default. However, to have Netcool/OMNIBus or Tivoli Enterprise Console events included in the Common Event Console, you must configure a connector for each of these event systems after you install the Tivoli Enterprise Portal. You might also want to change some of the configuration values for the default connector. For information on configuring the display of events from Netcool/OMNIBus and Tivoli Enterprise Console, see the *IBM Tivoli Monitoring Administrator's Guide*.

After situation forwarding is enabled, by default all situation events are forwarded to the specified event server. However, you can customize which situation events are forwarded and to which event server. For information on specifying which situation events to forward, see the Tivoli Enterprise Portal online help and the *IBM Tivoli Monitoring Tivoli Enterprise Portal User's Guide*.

Event integration scenarios: The scenarios in the next two chapters illustrate how event integration can be implemented with either Netcool/OMNIBus or Tivoli Enterprise Console by forwarding situation

events. To implement these scenarios, hub monitoring servers must be configured to forward events to at least one event server, and the event server must be configured to interpret the events and send updates to the originating monitoring server. A special event-synchronization component, the Situation Update Forwarder, must be installed on the event server. You can find complete instructions for implementing these scenarios in the next two chapters, Chapter 31, “Setting up event forwarding to Tivoli Enterprise Console,” on page 799 and Chapter 30, “Setting up event forwarding to Netcool/OMNIBus,” on page 697.

The System Monitor Agent, allows you to send OS monitoring data directly to Netcool/OMNIBus without first passing the data to a Tivoli Enterprise Monitoring Server. In this way these agents can run in agent-only environments that lack the standard Tivoli Monitoring servers (the Tivoli Enterprise Monitoring Server and the Tivoli Enterprise Portal Server). These monitoring agents, which run on Windows and on Linux/UNIX, effectively replace the OMNIBus System Service Monitors for monitoring of desktop operating systems. Chapter 12, “Monitoring your operating system via a System Monitor Agent,” on page 399 provides complete information on installing, configuring, and uninstalling a System Monitor Agent on either Windows or Linux.

Note: As of fix pack 2 for V6.2.2, for sites running x86_64 CPUs, both 32-bit and 64-bit Windows environments (Windows 2003, Vista, 2008) are supported for the System Monitor Agents. An enhancement provided with the first fix pack for version 6.2.2 enables the System Monitor Agents can send event data directly to OMNIBus, thus making a monitoring server unnecessary even for event processing.

- EIF events generated by autonomous agents (including the System Monitor Agents) can be sent directly to either Tivoli Enterprise Console or OMNIBus for private situations only.
- SNMP alerts generated by autonomous agents can be forwarded to any SNMP trap receiver, including OMNIBus's MTTRAPD probe for both enterprise and private situations.

For more information, see “Event forwarding from autonomous agents” on page 55.

For a complete list of operating systems that are supported for IBM Tivoli Monitoring components, see “Supported operating systems” on page 75. For event integration with Netcool/OMNIBus software prerequisites, see “Required software for event integration with Netcool/OMNIBus” on page 91.

Chapter 30. Setting up event forwarding to Netcool/OMNIBus

If you are already using Netcool/OMNIBus to monitor events from other sources in your enterprise, you can also view and manage situation events from a hub Tivoli Enterprise Monitoring Server or monitoring agent in the Netcool/OMNIBus Event List user interfaces.

Event integration requires Netcool/OMNIBus V7.2 or later and Netcool/OMNIBus Probe for Tivoli EIF version 10 or later. For all software prerequisites for event integration with Netcool/OMNIBus, see “Required software for event integration with Netcool/OMNIBus” on page 91.

Architecture overview

Event integration between a hub monitoring server (monitoring server) and Netcool/OMNIBus (event server) can be unidirectional or bidirectional.

By using the Tivoli Enterprise Portal or `tacmd` commands, you can create monitoring specifications called *situations* to detect when specific conditions or events occur in your environment and raise an event that is sent to Netcool/OMNIBus. Each situation is assigned (or distributed) to one or more managed systems that are to be monitored for a specific condition or a set of conditions.

Two types of events might be triggered by a situation: pure or sampled. When the determination of the event must be made based on observations made at specific intervals, the event is known as a *sampled event*. When the event is based on a spontaneous occurrence, the event is known as a *pure event*. Therefore, situations for sampled events have an interval associated with them, while those for pure events do not. In sampled events, the condition that caused the event can change, causing it to be no longer true. When a situation condition becomes true for a sampled situation, an event with open status is sent to Netcool/OMNIBus. When a situation condition is no longer true for a sampled situation, a status update event with the closed status is sent to Netcool/OMNIBus so that the event is cleared. Because pure events represent a spontaneous occurrence, an event with open status is sent to Netcool/OMNIBus each time the situation condition is true. No status update event is sent when the situation condition is not true. Therefore, pure events can be left open indefinitely unless you close them by using one of the following options:

- Close them by using the Tivoli Enterprise Portal Situation Event Console.
- Clear them in Netcool/OMNIBus.
- Define a situation Until modifier condition that automatically closes the pure event if a configured time interval passes or another situation condition becomes true.

Event integration between a hub monitoring server (monitoring server) and Netcool/OMNIBus (event server) can be unidirectional or bidirectional. Events can also be sent directly to Netcool/OMNIBus from IBM Tivoli Monitoring agents by using either SNMP or EIF. OMNIBus operators can use the Tivoli Netcool/OMNIBus WebGUI or native desktop environment to view events. By using the Netcool/OMNIBus Event List UI support in the WebGUI and native desktop, operators can acknowledge events, view event journals, take ownership of events, and run event management tools.

In a *unidirectional* architecture, hub monitoring servers use the Tivoli Event Integration Facility (EIF) interface to forward situation events to OMNIBus. The events are received by the Netcool/OMNIBus probe for Tivoli EIF, which maps them to OMNIBus events and then inserts them into the OMNIBus ObjectServer. If the status of a situation event changes, the hub monitoring server also sends status update events to Netcool/OMNIBus. The unidirectional architecture is similar when an agent is configured to send events to Netcool/OMNIBus. The situation events of the agent are forwarded to a

Netcool/OMNIbus probe for Tivoli EIF, or to the Netcool/OMNIbus SNMP probe that maps them to OMNIbus events. The agents also forward status update events to Netcool/OMNIbus when a sampled event condition is no longer true.

In a *bidirectional* architecture, when an OMNIbus operator acknowledges, deacknowledges, deletes, or clears a forwarded event, OMNIbus sends those status changes back to the hub monitoring server that forwarded them by using the IBM Tivoli Monitoring Situation Update Forwarder. (Severity changes, other than clearing an event, are not sent back to the hub monitoring server.) Bidirectional updates from Netcool/OMNIbus are supported only for events that originate from the hub monitoring server. Bidirectional updates are not supported for situation events sent directly from monitoring agents to SNMP, or EIF to Netcool/OMNIbus probes.

Use the bidirectional architecture in the following scenarios:

- If you want your Tivoli Enterprise Portal operators to see the same event status as your Netcool/OMNIbus operators and to acknowledge or deacknowledge events and to close pure events.
- If you want your Netcool/OMNIbus operators to be notified that an event condition has not been resolved when a sampled event is cleared or deleted by using the Netcool/OMNIbus Event List UI or event automations.

If you use the unidirectional architecture, the following conditions apply:

- Your Tivoli Portal Enterprise operators should not acknowledge, deacknowledge or close events.
- Pure events that are cleared in Netcool/OMNIbus remain open in the hub Tivoli Enterprise Monitoring Server. As a result, you should have a situation Until modifier configured for your pure events to close them automatically.
- Sampled events that are cleared in Netcool/OMNIbus are not re-opened in Netcool/OMNIbus until IBM Tivoli Monitoring detects that the situation condition has become false and then true again.

The event integration also supports single-tier, multitier and high availability Netcool/OMNIbus architectures. In a single-tier architecture without high availability, there is a single Netcool/OMNIbus ObjectServer. In a multitier architecture, there are multiple sets of ObjectServers for scalability purposes. You can add high availability to each of these architectures by adding a primary and backup ObjectServer to each tier. You can also configure EIF probes for peer-to-peer failover mode. For more information on configuring the event integration in multitier and high-availability architectures, see “Netcool/OMNIbus Multitiered and High-availability Architecture” on page 728.




Figure 145. A typical event flow showing both IBM Tivoli Monitoring and Netcool/OMNIBus environments for a single-tier architecture

The following steps outline the event flow for a typical bidirectional or unidirectional flow between IBM Tivoli Monitoring and Netcool/OMNIBus, with actions taken by Netcool/OMNIBus operators. Steps 1 to 5 are common to both bidirectional and unidirectional event flows. Steps 6 to 11 are specific to bidirectional event flows only.

1. IBM Tivoli Monitoring generates situation events that are sent to the Netcool/OMNIBus probe for Tivoli EIF. These events are also displayed in the Tivoli Enterprise Portal.
2. The probe maps a subset of the IBM Tivoli Monitoring EIF slots to the Netcool/OMNIBus ObjectServer attributes and creates an OMNIBus event.
3. The OMNIBus events are inserted into the Netcool/OMNIBus ObjectServer by using IBM Tivoli Monitoring provided triggers.
4. Events in the Netcool/OMNIBus ObjectServer are displayed in the Netcool/OMNIBus Native Event List or WebGUI and in the Tivoli Enterprise Portal.
5. Forwarded events are acknowledged, deacknowledged, deleted, or cleared by Netcool/OMNIBus operators or by automation within OMNIBus or Impact. Event status can also be updated by an integrated trouble ticket system or by integration with another Event Server.
6. IBM Tivoli Monitoring triggers in the Netcool/OMNIBus ObjectServer database forward the event status changes to the IBM Tivoli Monitoring Situation Update Forwarder.
7. The IBM Tivoli Monitoring Situation Update Forwarder sends SOAP requests to the hub monitoring server.

8. Status changes are propagated through the Tivoli Enterprise Portal Server and shown in the Tivoli Enterprise Portal. The complete list of open, acknowledged, and deacknowledged events is shown in the Situation Event Console workspace of the Tivoli Enterprise Portal.
9. The event status change is also sent back to Netcool/OMNIBus and to any other event destinations that are configured for the situation.
10. The probe maps a subset of the IBM Tivoli Monitoring EIF slots to the Netcool/OMNIBus ObjectServer attributes and creates an OMNIBus event.
11. IBM Tivoli Monitoring triggers determine the event is a loopback event and ignores it.

A Tivoli Enterprise Portal operator can also acknowledge and deacknowledge events or close pure events for the bidirectional architecture. The Tivoli Enterprise Portal Server processes the event status changes from the Tivoli Enterprise Portal and notifies the hub Tivoli Enterprise Monitoring Server of the changes. The event status updates are sent by the hub Tivoli Enterprise Monitoring Server to Netcool/OMNIBus by using the flows shown in Figure 145 on page 699.

Event behavior

How events are handled depends on several criteria, including the architecture type (unidirectional versus bidirectional) and the event type (pure versus sampled).

Table 125 describes the behavior for events sent from the Hub monitoring server to Netcool/OMNIBus. Table 126 on page 712 describes the behavior for events sent directly from the monitoring agents to Netcool/OMNIBus.

Table 125. Behavior of events originating from a hub monitoring server

Action	Event type	Unidirectional behavior	Bidirectional behavior
Situation condition becomes true.	Pure and sampled events.	<p>Summary: A new event is opened in the hub monitoring server and in the Netcool/OMNIBus ObjectServer if it does not deduplicate an existing event.</p> <p>Details: The hub monitoring server opens a new situation event and sends an event with the Open status to Netcool/OMNIBus ObjectServer using flows 1 to 4 as shown in Figure 145 on page 699 and a new event is opened in Netcool/OMNIBus if it does not deduplicate an existing event.</p> <p>Note: The hub monitoring server sends an event with open status to Netcool/OMNIBus each time a condition is true for a pure situation. For sampled situations, an event with open status is sent when the situation condition transitions from not true to true. Until the sampled situation condition becomes false, another event is not sent unless the status of the event is changed, for example to acknowledged.</p>	Same as unidirectional behavior.

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
Sampled event situation condition is no longer true or a pure sampled situation UNTIL modifier condition becomes true.	Pure and sampled events.	<p>Summary: Event is closed in the hub monitoring server and cleared in the Netcool/OMNIBus ObjectServer.</p> <p>Details: After the event is closed in the hub monitoring server, a closed status update event is sent to Netcool/OMNIBus ObjectServer by using flows 1 to 4 as shown in Figure 145 on page 699. When the IBM Tivoli Monitoring triggers process the status update event they clear the event in the Netcool/OMNIBus ObjectServer.</p> <p>Note: For more information about creating a situation UNTIL modifier and details on when sampled events are closed if they have an UNTIL modifier, see the <i>IBM Tivoli Monitoring Tivoli Enterprise Portal User's Guide</i>.</p>	Same as unidirectional behavior.
Event acknowledged using the Netcool/OMNIBus Event List UI.	Pure and sampled events.	<p>Summary: Event status is changed to acknowledged in the Netcool/OMNIBus ObjectServer. However, the event status is not updated in the hub monitoring server and Tivoli Enterprise Portal.</p>	<p>Summary: The event status is changed to acknowledged in the Netcool/OMNIBus ObjectServer, the hub monitoring server, and Tivoli Enterprise Portal.</p> <p>Details: After the event status is changed to acknowledged in the Netcool/OMNIBus ObjectServer, the hub monitoring server is notified about the status change by flows 5 through 8 in Figure 145 on page 699. In flow 7, the IBM Tivoli Monitoring Situation Update Forwarder sends a CT_Acknowledge SOAP request to the hub monitoring server. The hub monitoring server changes the event status to Acknowledged when it processes the SOAP request and sends a status update event back to OMNIBus using flows 9 through 11.</p>
Event cleared or deleted using the Netcool/OMNIBus Event List UI.	Pure events.	<p>Summary: The pure event is cleared or deleted in the Netcool/OMNIBus ObjectServer. However, the event remains open in the hub monitoring server and Tivoli Enterprise Portal until the UNTIL modifier condition of the situation becomes true.</p>	<p>Summary: The pure event is cleared or deleted in the Netcool/OMNIBus ObjectServer and closed in the hub monitoring server and Tivoli Enterprise Portal.</p> <p>Details: After the event is cleared or deleted in the Netcool/OMNIBus ObjectServer, the hub monitoring server is notified about the status change by flows 5 through 8 in Figure 145 on page 699. In flow 7, the IBM Tivoli Monitoring Situation Update Forwarder sends a CT_Reset SOAP request to the hub monitoring server. The hub monitoring server closes the event when it processes the SOAP request and sends a status update event back to OMNIBus using flows 9 through 11.</p>

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
Event cleared or deleted using the Netcool/OMNIBus Event List UI.	Sampled events.	<p>Summary: The sampled event is cleared or deleted in the Netcool/OMNIBus ObjectServer. However, the event remains open in the hub monitoring server and Tivoli Enterprise Portal until the situation condition is no longer true. No further status updates are sent to Netcool/OMNIBus until the situation condition becomes false and then true again. Therefore, the Netcool/OMNIBus operator is not notified that the event condition has not been resolved.</p>	<p>Summary: The sampled event is cleared or deleted in the Netcool/OMNIBus ObjectServer but its status is changed to <i>Acknowledged</i> in the hub monitoring server and Tivoli Enterprise Portal for a specified time. If the situation condition is still true after the specified time, a status update event is sent to Netcool/OMNIBus and an event is opened. This status update notifies the Netcool/OMNIBus operator that the event condition is not resolved.</p> <p>Details: When the sampled event is cleared or deleted, the event data is cached by the ObjectServer in an IBM Tivoli Monitoring table. Then the hub monitoring server is notified of the status change by flows 5 through 8 in Figure 145 on page 699. In flow 7, the IBM Tivoli Monitoring Situation Update Forwarder sends a CT_Acknowledge SOAP request with a configurable timeout to the hub monitoring server. The hub monitoring server changes the event status to <i>Acknowledged</i> and starts an expiration timer. A status update event is sent back to OMNIBus by using flows 9 through 11. The events are marked as <i>Acknowledged</i> in the hub monitoring server and Tivoli Enterprise Portal because a sampled event cannot be closed unless the situation condition is no longer true. By leaving the situation as <i>Acknowledged</i>, Netcool/OMNIBus is notified if the situation condition is still true after the timeout expires. If the situation condition is still true when the timeout expires, the hub monitoring server sends an <i>Acknowledgement Expired</i> status update event to Netcool/OMNIBus ObjectServer by using flows 1 to 4 as shown in Figure 145 on page 699. If the event has already been removed from the Netcool/OMNIBus ObjectServer alerts.status table, a new event is opened in the ObjectServer.</p>
Event cleared or deleted using the Netcool/OMNIBus Event List UI (continued)			<p>Details: (continued) Because status update events contain only base ITM EIF slots and no agent-specific slots, the event is re-opened using data that was cached when the event was cleared or deleted from the Netcool/OMNIBus Event List UI. However, if the event is still in the Netcool/OMNIBus ObjectServer alerts.status table when the status update event is processed, the event is deduplicated by the ITM triggers. The event is then reopened and contains event attribute settings from the original event.</p>

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
Event deacknowledged by using the Netcool/OMNIBus Event List UI.	Pure and sampled events.	Summary: The event status is changed to deacknowledged in the Netcool/OMNIBus ObjectServer. However, the event status is not updated in the hub monitoring server and Tivoli Enterprise Portal.	Summary: The event status is changed to deacknowledged in the Netcool/OMNIBus ObjectServer and to resurfaced in the hub monitoring server and Tivoli Enterprise Portal. Details: After the event status is changed to deacknowledged in the Netcool/OMNIBus ObjectServer, the hub monitoring server is notified about the status change by flows 5 through 8 in Figure 145 on page 699. In flow 7, the IBM Tivoli Monitoring Situation Update Forwarder sends a CT_Resurface request to the hub monitoring server. The hub monitoring server changes the event status to Resurfaced when it processes the SOAP request and sends a status update event back to OMNIBus using flows 9 through 11.
Event acknowledged without a timeout by using the Tivoli Enterprise Portal Situation Event Console.	Pure and sampled events.	Tivoli Enterprise Portal operators should not change the event status when the unidirectional architecture is being used.	Summary: The event status is changed to acknowledged in both the hub monitoring server and Netcool/OMNIBus ObjectServer. Details: After the event status is changed to acknowledged in the hub monitoring server, an Acknowledged status update event is sent to Netcool/OMNIBus using flows 1 to 4 as shown in Figure 145 on page 699 and the event status is changed to Acknowledged in the Netcool/OMNIBus ObjectServer and Netcool/OMNIBus Event List UI.

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
Event acknowledged with timeout by using the Tivoli Enterprise Portal Situation Event Console.	Pure and sampled events.	Tivoli Enterprise Portal operators should not change the event status when the unidirectional architecture is being used.	<p>Summary: The event status is changed to acknowledged in both the hub monitoring server and Netcool/OMNIbus ObjectServer and you can configure how Netcool/OMNIbus handles the timeout notification event. See the detailed description for more information.</p> <p>Details: After the event status is changed to acknowledged in the hub monitoring server, an Acknowledged status update event is sent to Netcool/OMNIbus using flows 1 to 4 as shown in Figure 145 on page 699 and the event status is changed to Acknowledged in the Netcool/OMNIbus ObjectServer and Netcool/OMNIbus Event List UI. When the timeout expires and the situation event is still open in the hub monitoring server, the hub monitoring server sets the event status to Acknowledgement Expired and sends an Acknowledgement Expired status update event to Netcool/OMNIbus using flows 1 to 4 as shown in Figure 145 on page 699.</p> <p>By default, the IBM Tivoli Monitoring triggers reject the Acknowledgement Expired status update event and use flows 5 to 8 to send a request to the hub monitoring server to set the event status to Acknowledged. (In flow 7, the Situation Update Forwarder sends a CT_Acknowledge request to the hub monitoring server.) The hub monitoring server sets the event status to Acknowledged when it processes the SOAP request and sends a status update event back to OMNIbus using flows 9 through 11.</p> <p>You can override the default IBM Tivoli Monitoring trigger behavior by setting the <i>sit_ack_expired_def_action</i> variable to ACCEPT using the procedure described in "Customizing how the IBM Tivoli Monitoring OMNIbus triggers handle event status updates from the monitoring servers" on page 775. If you set the variable to ACCEPT, the IBM Tivoli Monitoring triggers deacknowledge the event in the Netcool/OMNIbus ObjectServer but the event still has the acknowledgment expired status in the hub monitoring server.</p>

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
Event deacknowledged using the Tivoli Enterprise Portal Situation Event Console.	Pure and sampled events.	Tivoli Enterprise Portal operators should not change the event status when the unidirectional architecture is being used.	<p>Summary: Behavior is configurable. See the detailed description for the two types of behaviors that are supported.</p> <p>Details: After the event status is changed to Resurfaced in the hub monitoring server, a Resurfaced status update event is sent to Netcool/OMNIBus using flows 1 to 4 as shown in Figure 145 on page 699. By default the IBM Tivoli Monitoring triggers accept the Resurfaced status update event and deacknowledge the event in the Netcool/OMNIBus ObjectServer.</p> <p>You can override the default trigger behavior by setting the <code>sit_resurface_def_action</code> variable to REJECT using the procedure described in “Customizing how the IBM Tivoli Monitoring OMNIBus triggers handle event status updates from the monitoring servers” on page 775. If you set the variable to REJECT then the IBM Tivoli Monitoring triggers use flows 5 to 8 in Figure 145 on page 699 to send a CT_ACKNOWLEDGE SOAP request to the hub monitoring server. The hub monitoring server sets the event status to Acknowledged when it processes the SOAP request and sends a status update event back to OMNIBus using flows 9 through 11.</p>

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
<p>Situation is stopped.</p> <p>Note: Situations are stopped if an operator initiates the situation stop action from the Tivoli Enterprise Portal or modifies the situation definition. However, a situation is not stopped if its distribution list is modified.</p>	Pure events.	<p>Summary: The hub monitoring server closes all pure events for the situation. These same events are cleared in the Netcool/OMNIBus ObjectServer unless you configure a different behavior in the IBM Tivoli Monitoring triggers.</p> <p>Details: As the hub monitoring server closes all pure events for the situation, it sends a situation stop event to Netcool/OMNIBus for each remote monitoring server that was monitoring the situation. The situation stop event specifies the remote monitoring server in the situation_thruNode EIF slot. (This slot is mapped to the ITMThruNodeOMNIBus attribute by the Tivoli Netcool/OMNIBus EIF Probe.)</p> <p>When the IBM Tivoli Monitoring triggers in Netcool/OMNIBus process a situation stop event, they clear all events for the situation that are detected by the remote monitoring server specified by the ITMThruNode OMNIBus attribute. However, you can configure the IBM Tivoli Monitoring triggers to ignore situation stop events for pure events. For more information, see “Customizing how the IBM Tivoli Monitoring OMNIBus triggers handle event status updates from the monitoring servers” on page 775.</p>	Same as unidirectional behavior.
<p>Situation is stopped.</p> <p>Note: Situations are stopped if an operator initiates the situation stop action from the Tivoli Enterprise Portal or modifies the situation definition. However, a situation is not stopped if its distribution list is modified.</p>	Sampled events.	<p>Summary: The hub monitoring server closes all sampled events for the situation. These same events are cleared in the Netcool/OMNIBus ObjectServer.</p> <p>Details: As the hub monitoring server closes all sampled events for the situation, it sends a situation stop event to Netcool/OMNIBus for each remote monitoring server that was monitoring the situation. The situation stop event specifies the remote monitoring server in the situation_thruNode EIF slot. (This slot is mapped to the ITMThruNode OMNIBus attribute by the Tivoli Netcool/OMNIBus EIF Probe.)</p> <p>When the IBM Tivoli Monitoring triggers in Netcool/OMNIBus process a situation stop event, they clear all events for the situation that are detected by the remote monitoring server specified by the ITMThruNode OMNIBus attribute.</p>	Same as unidirectional behavior.

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
Agent stopped.	Pure events.	<p>Summary: Stopping the agent has no effect on pure situation event status. An MS_Offline situation event is also sent to Netcool/OMNIBus to indicate that the monitoring agent is not being monitored.</p> <p>Details: After the monitoring server of the agent detects that the agent has not responded for three situation sampling intervals, the situation's event is closed in the hub monitoring server. A Closed status update event is sent to Netcool/OMNIBus ObjectServer by using flows 1 to 4 as shown in Figure 145 on page 699. When the IBM Tivoli Monitoring triggers process the status update event, they clear the event in the Netcool/OMNIBus ObjectServer.</p> <p>If you do not want events to be closed in Netcool/OMNIBus after an agent is stopped, you can customize this behavior. For more information, see "Customizing event status processing behavior when agent switching is used or the agent goes offline" on page 782.</p>	Same as unidirectional behavior.
Agent stopped.	Sampled events.	<p>Summary: The sampled events from the agent are closed in the hub monitoring server if the agent is stopped for long enough. Its events are also cleared in the Netcool/OMNIBus ObjectServer. An MS_Offline situation event is also sent to Netcool/OMNIBus to indicate that the monitoring agent is not being monitored.</p> <p>Details: After the monitoring server of the agent detects that the agent has not responded for three situation sampling intervals, the situation's event is closed in the hub monitoring server. A Closed status update event is sent to Netcool/OMNIBus ObjectServer by using flows 1 to 4 as shown in Figure 145 on page 699. When the IBM Tivoli Monitoring triggers process the status update event, the triggers clear the event in the Netcool/OMNIBus ObjectServer.</p> <p>If you do not want events to be closed in Netcool/OMNIBus after an agent is stopped, you can customize this behavior. For more information, see "Customizing event status processing behavior when agent switching is used or the agent goes offline" on page 782.</p>	Same as unidirectional behavior.

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
Agent loses connectivity to the primary monitoring server and switches to the secondary monitoring server.	Pure and sampled events.	<p>When an agent has a situation that has triggered and the situation is true, and the agent subsequently loses connectivity to the Tivoli Enterprise Monitoring Server and switches to a different monitoring server, the Hub monitoring server sends an EIF event with class type ITM_ControlSignal to the Netcool/OMNIBus ObjectServer for each situation event that is still open for the agent. When the IBM Tivoli Monitoring triggers process this event, they update the ITMThruNode attribute to specify the new monitoring server for the agent. However, the event might be closed by the original monitoring server if it detects the agent is no longer responding to it before the new monitoring server determines that the situation event is still true and the Hub monitoring server sends the ITM_ControlSignal event.</p> <p>There are environment variables that can be added to the monitoring server environment file to customize the behavior of event status processing when agent switching occurs to help ensure that events are not closed by the original monitoring server. For a complete description of these variables, see “Customizing event status processing behavior when agent switching is used or the agent goes offline” on page 782.</p> <p>Note: If an agent switches to another monitoring server because the original monitoring server was stopped, all sampled events for the agent will be closed by the original monitoring server. This behavior is not configurable. However, you can configure whether pure events are closed in this scenario. For more information, see “Customizing how the IBM Tivoli Monitoring OMNIBus triggers handle event status updates from the monitoring servers” on page 775.</p>	Same as unidirectional behavior.
Hub monitoring server stopped.	Pure and sampled events.	<p>Summary: No flows occur between the hub monitoring server and Netcool/OMNIBus when the hub monitoring server is stopped.</p>	Same as unidirectional behavior.

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
Hub monitoring server started.	Pure events.	<p>Summary: The pure events in the Netcool/OMNIbus ObjectServer are unaffected. However, the hub monitoring server and Netcool/OMNIbus might not have the same status for pure events.</p> <p>Details: When the hub monitoring server is started, it sends a master_reset event to Netcool/OMNIbus using flows 1 to 4 as shown in Figure 145 on page 699. The IBM Tivoli Monitoring triggers in the Netcool/OMNIbus ObjectServer do not update the status of pure events when the master reset event is processed.</p> <p>However, Netcool/OMNIbus and the hub monitoring server (and Tivoli Enterprise Portal) might have a different status for the pure events after the hub is restarted.</p> <ul style="list-style-type: none"> • The hub monitoring server does not have any event status for pure events that were opened or acknowledged prior to the hub restart if the events were for agents connected directly to the hub monitoring server. These events might still be open or acknowledged in the Netcool/OMNIbus ObjectServer. • The hub monitoring server has a status of open for pure events that were open or acknowledged prior to the hub restart if the events were for agents connected to the remote Tivoli Enterprise Monitoring Server. However, these events may be acknowledged, cleared, or deleted in Netcool/OMNIbus. 	Same as unidirectional behavior.

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
Hub monitoring server started.	Sampled events.	<p>Summary: The sampled events in the Netcool/OMNIBus ObjectServer from this hub monitoring event are cleared.</p> <p>Details: When the hub monitoring server is started, it sends a master_reset event to Netcool/OMNIBus using flows 1 to 4 as shown in Figure 145 on page 699. The IBM Tivoli Monitoring triggers in the Netcool/OMNIBus ObjectServer clear all sampled events from this hub monitoring server when the master reset event is processed. The master reset handling ensures that events are cleared in Netcool/OMNIBus if the situation condition became false while the hub monitoring server was stopped. The events whose situation conditions are still true will be reopened in Netcool/OMNIBus after the master reset event is sent.</p>	Same as unidirectional behavior.

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
Remote monitoring server stopped.	Pure events.	<p>Summary: The hub monitoring server closes all of the pure events for agents connected to the remote monitoring server. The same events are cleared in Netcool/OMNIbus Object unless you configure a different behavior. An MS_Offline situation event is sent to Netcool/OMNIbus for the remote monitoring server to indicate that these managed systems are not being monitored. You will not see an MS_Offline message for each of the monitoring agents when the remote monitoring server goes offline.</p> <p>Details: As the hub monitoring server closes the pure events for each situation being monitored by the remote monitoring server, it sends a situation stop event to Netcool/OMNIbus and specifies the remote monitoring server in the situation_thruNode EIF slot. (This slot is mapped to the ITMThruNode OMNIbus attribute by the Tivoli Netcool/OMNIbus EIF Probe.)</p> <p>When the IBM Tivoli Monitoring triggers in Netcool/OMNIbus process a situation stop event, they clear all events for the situation that are detected by the remote monitoring server specified by the ITMThruNode OMNIbus attribute. However, you can configure the IBM Tivoli Monitoring triggers to ignore situation stop events for pure events. For more information, see “Customizing how the IBM Tivoli Monitoring OMNIbus triggers handle event status updates from the monitoring servers” on page 775.</p>	Same as unidirectional behavior.

Table 125. Behavior of events originating from a hub monitoring server (continued)

Action	Event type	Unidirectional behavior	Bidirectional behavior
Remote monitoring server stopped.	Sampled events.	<p>Summary: The hub monitoring server closes all of the sampled events for agents connected to the remote monitoring server. The same events are cleared in Netcool/OMNIbus Object. An MS_Offline situation event is sent to Netcool/OMNIbus for the remote monitoring server and each of the monitoring agents connected to the monitoring server to indicate that these managed systems are not being monitored. You will not see an MS_Offline message for each of the monitoring agents when the remote monitoring server goes offline.</p> <p>Details: As the hub monitoring server closes the sampled events for each situation being monitored by the remote monitoring server, it sends a situation stop event to Netcool/OMNIbus and specifies the remote Tivoli Enterprise Monitoring Server in the situation_thrnode EIF slot. (This slot is mapped to the ITMThruNode OMNIbus attribute by the Tivoli Netcool/OMNIbus EIF Probe.)</p> <p>When the IBM Tivoli Monitoring triggers in Netcool/OMNIbus process a situation stop event, they clear all events for the situation that are detected by the remote monitoring server specified by the ITMThruNode OMNIbus attribute.</p>	Same as unidirectional behavior.
Remote monitoring server started.	Pure and sampled events.	Same as unidirectional behavior when a remote monitoring server is stopped.	Same as bidirectional behavior when a remote monitoring server is stopped.

Table 126. Behavior of events originating from IBM Tivoli Monitoring agents

Action	Event Type	Behavior
Situation condition becomes true.	Pure and sampled events.	<p>A new event is opened in the agent and in the Netcool/OMNIbus ObjectServer if it does not deduplicate an existing event.</p> <p>An event with open status is sent from the agent to Netcool/OMNIbus each time the situation condition is true, if the following conditions are met:</p> <ul style="list-style-type: none"> • The situation generates pure events, or • The situation generates sampled events, the events are sent to Netcool/OMNIbus using SNMP, and the situation mode is set to Rising Continuous.
Sampled event situation condition is no longer true.	Sampled events.	Event is closed in the agent and cleared in the Netcool/OMNIbus ObjectServer.

Table 126. Behavior of events originating from IBM Tivoli Monitoring agents (continued)

Action	Event Type	Behavior
Event acknowledged or deacknowledged using Netcool/OMNIBus Event List UI in OMNIBus.	Pure and sampled events.	The event status is changed to acknowledged or deacknowledged in Netcool/OMNIBus but the event status maintained by the agent is unaffected.
Event cleared or deleted using Netcool/OMNIBus Event List UI.	Pure events.	The event is cleared or deleted in the Netcool/OMNIBus ObjectServer. However, the event status maintained by the agent is not affected.
Event cleared or deleted using Netcool/OMNIBus Event List UI.	Sampled events.	<p>The event is cleared or deleted in the Netcool/OMNIBus ObjectServer. However, the event status maintained by the agent is unaffected and the Netcool/OMNIBus operator is not notified if the event condition has not been resolved, unless the following conditions are met:</p> <ul style="list-style-type: none"> • The agent has been configured to send SNMP events to Netcool/OMNIBus, and • The situation mode is set to Rising Continuous so that an event is sent to Netcool/OMNIBus each sampling interval that the situation event evaluates to true. With this mode, the event is reopened in Netcool/OMNIBus if the event condition is still true.
Situation is stopped using the Agent Service Interface.	Pure and sampled events.	<p>If the agent is configured to send lifecycle events when a situation is stopped, a EE_SIT_STOPPED event is sent to Netcool/OMNIBus.</p> <p>For agents that send SNMP events to Netcool/OMNIBus, the events in Netcool/OMNIBus are not affected by this lifecycle event.</p> <p>For agents that send EIF events to Netcool/OMNIBus, the events from the agent for the stopped situation are cleared in Netcool/OMNIBus.</p>
Agent is stopped.	Pure and sampled events.	No events are sent by the agent when it is stopped.

Table 126. Behavior of events originating from IBM Tivoli Monitoring agents (continued)

Action	Event Type	Behavior
Agent is started.	Pure and sampled events.	<p>If the agent is configured to send SNMP events to Netcool/OMNIBus, no events other than lifecycle traps are sent by the agent when it is started.</p> <p>If the agent is configured to send EIF events to Netcool/OMNIBus, by default the agent does not send any events (other than lifecycle events) to Netcool/OMNIBus when it is started. However, you can change this behavior and configure the agent to send a master reset event to Netcool/OMNIBus when the agent is started. When the IBM Tivoli Monitoring triggers in Netcool/OMNIBus process this event, they clear all events for the agent. This ensures that events are cleared in Netcool/OMNIBus if the situation condition became false while the agent was stopped. The events whose situation conditions are still true will be reopened in Netcool/OMNIBus after the master reset event is sent.</p> <p>Note: The agent does not maintain event status across restarts.</p>

Agents can also send lifecycle and heartbeat events to Netcool/OMNIBus. For more details on lifecycle and heartbeat events, see the Agent Autonomy chapter in the *IBM Tivoli Monitoring Administrator's Guide*.

IBM Tivoli Monitoring Event Synchronization component

You must install the IBM Tivoli Monitoring Event Synchronization component regardless of the architectural solution you choose for event integration between the hub monitoring server and Netcool/OMNIBus.

This component consists of the following three parts:

IBM Tivoli Monitoring rules file for the Netcool/OMNIBus Probe for Tivoli EIF.

Used to update the probe so it can understand IBM Tivoli Monitoring situation events.

SQL files to update the Netcool/OMNIBus ObjectServer database.

Used to update the Netcool/OMNIBus ObjectServer database schema that contain attributes specific to Tivoli Monitoring, and create or update the triggers that process IBM Tivoli Monitoring events. The SQL files also forward events to the IBM Tivoli Monitoring Situation Update Forwarder when using bidirectional communication.

IBM Tivoli Monitoring Situation Update Forwarder executable.

The program that forwards the updates from OMNIBus back to IBM Tivoli Monitoring.

The IBM Tivoli Monitoring rules file for the probe maps a subset of the IBM Tivoli Monitoring EIF slots to OMNIBus ObjectServer attributes to create an OMNIBus event. The SQL files update the Netcool/OMNIBus ObjectServer database schema to contain IBM Tivoli Monitoring specific attributes and create or update the triggers that process IBM Tivoli Monitoring events as well as forward events to the IBM Tivoli Monitoring Situation Update Forwarder when using bidirectional communication. The IBM Tivoli Monitoring Situation Update Forwarder is used to forward updates to the situation events, back to the originating hub monitoring server via SOAP messages.

If you choose to send EIF events from agents to Netcool/OMNIBus, you must also install the IBM Tivoli Monitoring Event Synchronization component because it contains the IBM Tivoli Monitoring rules file for the Netcool/OMNIBus Probe for EIF, and the SQL files to update the Netcool/OMNIBus ObjectServer database.

If you are sending only SNMP events from agents to Netcool/OMNIBus, you do not need to install the synchronization component. The Tivoli Monitoring rules file for the Netcool/OMNIBus Probe for SNMP, and the SQL files to update the Netcool/OMNIBus ObjectServer database, are packaged separately.

If IBM Tivoli Monitoring situations have custom slots, you must update OMNIBus to add those custom slots. For more information about making this change, see “Configuring the Netcool/OMNIBus EIF probe” on page 760.

Architecture scenarios

The scenarios in this section illustrate ways to forward situation events from one or more Tivoli Enterprise Monitoring Servers to one or more Netcool/OMNIBus ObjectServers.

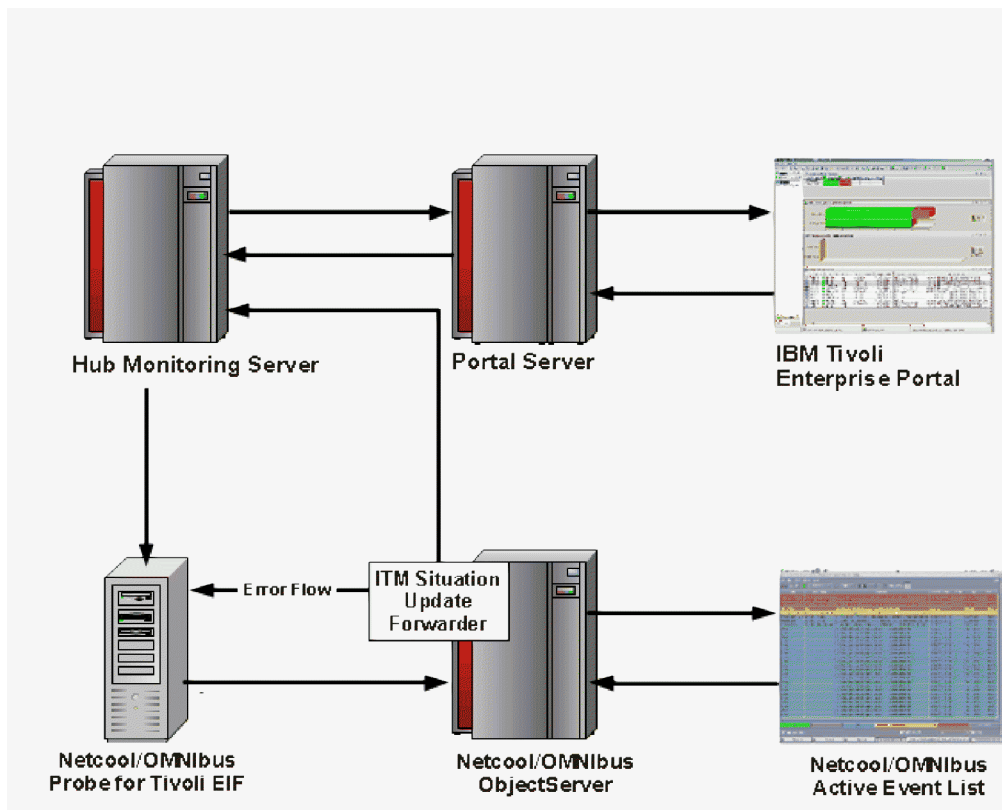
These scenarios use the bidirectional architecture, as described in the previous sections. If unidirectional architecture is applied to these scenarios, the IBM Tivoli Monitoring Situation Update Forwarder is removed and Netcool/OMNIBus does not send event status updates back to IBM Tivoli Monitoring.

One hub Tivoli Enterprise Monitoring Server and one Netcool/OMNIBus ObjectServer

In this scenario, a single hub Tivoli Enterprise Monitoring Server is forwarding events to a single Netcool/OMNIBus ObjectServer.

With bidirectional architecture, as changes are made to those events in the Netcool/OMNIBus ObjectServer, the updates are forwarded back to the hub monitoring server.

By default, event forwarding is not enabled for a new situation unless you base the new situation definition on an existing situation that already has event forwarding enabled or you explicitly enable event forwarding. When you enable event forwarding for a situation, the hub monitoring server sends the situation events to the EIF probe that was specified when event forwarding was enabled for the hub monitoring server.



Uses

This scenario describes the most basic of event synchronization architectures. This type of architecture is useful for small environments. Other uses include proof of concept or test environments. You can install all components on the same server. However, for performance reasons, install components on separate servers when setting up production environments.

Installation and configuration

If you are installing IBM Tivoli Monitoring and Netcool/OMNIBus event integration for the first time, complete the tasks in this table.

However, if you are upgrading an existing event integration environment, see “Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIbus integration” on page 766.

Table 127. Installation and configuration: one hub Tivoli Enterprise Monitoring Server and one Netcool/OMNIBus ObjectServer

	Task	Architecture Type	Administrator
1.	<p>Install Netcool/OMNIBus ObjectServer if it is not already installed or upgrade it to the fix pack version required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIBus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIBus” on page 91.</p>	Unidirectional and bidirectional	Netcool/OMNIBus
2.	<p>Install the Netcool/OMNIBus Probe for Tivoli EIF if it is not already installed or upgrade it to the release required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIBus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIBus” on page 91.</p>	Unidirectional and bidirectional	Netcool/OMNIBus
3.	“Installing the IBM Tivoli Monitoring Event Synchronization Component” on page 739.	Unidirectional and bidirectional	Netcool/OMNIBus
4.	“Updating the OMNIBus database schema on single-tier or aggregation tier ObjectServers” on page 753.	Unidirectional and bidirectional	Netcool/OMNIBus
5.	“Changing the default deduplication trigger” on page 755.	Unidirectional and bidirectional	Netcool/OMNIBus
6.	“Configuring the OMNIBus server for program execution from scripts” on page 763.	Bidirectional	Netcool/OMNIBus
7.	Start the IBM Tivoli Monitoring Situation Update Forwarder. For more information, see “Starting and stopping the IBM Tivoli Monitoring Situation Update Forwarder” on page 759.	Bidirectional	Netcool/OMNIBus
8.	“Configuring the Netcool/OMNIBus EIF probe” on page 760.	Unidirectional and bidirectional	Netcool/OMNIBus

Table 127. Installation and configuration: one hub Tivoli Enterprise Monitoring Server and one Netcool/OMNIBus ObjectServer (continued)

	Task	Architecture Type	Administrator
9.	"Configuring the hub monitoring server to forward events" on page 764.	Unidirectional and bidirectional	IBM Tivoli Monitoring
10.	"Verifying installation and configuration" on page 765.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIBus
11.	Determine if additional configuration tasks should be performed. See "Customizing Event Integration" on page 772.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIBus

Multiple hub Tivoli Enterprise monitoring servers and one Netcool/OMNIBus ObjectServer

In this scenario, multiple hub Tivoli Enterprise Monitoring Servers are forwarding events to a single Netcool/OMNIBus ObjectServer.

As changes are made to those events in the Netcool/OMNIBus ObjectServer, the updates are forwarded back to the hub monitoring server that is associated with that situation. Event forwarding must be enabled on each monitoring server, and the EIF probe associated with the ObjectServer must be defined as the default EIF receiver for each hub monitoring server.

By default, event forwarding is not enabled for a new situation unless you base the new situation definition on an existing situation that already has event forwarding enabled or you explicitly enable event forwarding. When you enable event forwarding for a situation, the hub monitoring server sends the situation's events to the EIF probe that was specified when event forwarding was enabled for the hub monitoring server.




Figure 147. Multiple hub Tivoli Enterprise Monitoring Servers and one Netcool/OMNIBus ObjectServer

Uses

This scenario demonstrates how you can consolidate all your IBM Tivoli Monitoring events by forwarding them to the same Netcool/OMNIBus ObjectServer through the Netcool/OMNIBus Probe for Tivoli EIF.

Installation and configuration

If you are installing IBM Tivoli Monitoring and Netcool/OMNIBus event integration for the first time, perform the tasks in this table.

However, if you are upgrading an existing event integration environment, see “Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIBus integration” on page 766.

Table 128. Installation and configuration: multiple hub Tivoli Enterprise Monitoring Servers and one Netcool/OMNIbus ObjectServer

	Task	Architecture Type	Administrator
1.	<p>Install Netcool/OMNIbus ObjectServer if it is not already installed or upgrade it to the fix pack version required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIbus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIbus” on page 91.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
2.	<p>Install the Netcool/OMNIbus Probe for Tivoli EIF if it is not already installed or upgrade it to the release required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIbus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIbus” on page 91.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
3.	“Installing the IBM Tivoli Monitoring Event Synchronization Component” on page 739.	Unidirectional and bidirectional	Netcool/OMNIbus
4.	“Updating the OMNIbus database schema on single-tier or aggregation tier ObjectServers” on page 753.	Unidirectional and bidirectional	Netcool/OMNIbus
5.	“Changing the default deduplication trigger” on page 755.	Unidirectional and bidirectional	Netcool/OMNIbus
6.	“Configuring the OMNIbus server for program execution from scripts” on page 763.	Bidirectional	Netcool/OMNIbus
7.	Start the IBM Tivoli Monitoring Situation Update Forwarder. For more information, see “Starting and stopping the IBM Tivoli Monitoring Situation Update Forwarder” on page 759.	Bidirectional	Netcool/OMNIbus

Table 128. Installation and configuration: multiple hub Tivoli Enterprise Monitoring Servers and one Netcool/OMNIBus ObjectServer (continued)

	Task	Architecture Type	Administrator
8.	Define additional monitoring servers to the IBM Tivoli Monitoring Situation Update Forwarder. For more information, see “Updating the IBM Tivoli Monitoring Situation Forwarder to forward event status updates to additional monitoring servers” on page 774.	Bidirectional	Netcool/OMNIBus
9.	“Configuring the Netcool/OMNIBus EIF probe” on page 760.	Unidirectional and bidirectional	Netcool/OMNIBus
10.	<i>For each hub monitoring server:</i> “Configuring the hub monitoring server to forward events” on page 764.	Unidirectional and bidirectional	IBM Tivoli Monitoring
11.	“Verifying installation and configuration” on page 765.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIBus
12.	Determine if additional configuration tasks should be performed. See “Customizing Event Integration” on page 772.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIBus

One hub Tivoli Enterprise Monitoring Server and multiple Netcool/OMNIBus ObjectServers

In this scenario, one hub Tivoli Enterprise Monitoring Server is forwarding events to multiple Netcool/OMNIBus ObjectServers where each ObjectServer is configured for the Netcool/OMNIBus single-tier architecture.

Each event server to which events are forwarded must have an associated EIF probe and an installed IBM Tivoli Monitoring Situation Update Forwarder. For each situation, you must specify the event server to which the situation event should be forwarded.

By default, event forwarding is not enabled for a new situation unless you base the new situation definition on an existing situation that already has event forwarding enabled, or you explicitly enable event forwarding for the situation. When you enable event forwarding for a situation, by default the hub monitoring server sends the situation’s events to the EIF probe that was specified when event forwarding was enabled for the hub monitoring server. If you want to forward events to multiple Netcool/OMNIBus ObjectServers, you must use the `tacmd createEventDest` command to create additional event server destination definitions and then select the appropriate event destination for each situation that has event forwarding enabled.




Figure 148. One hub Tivoli Enterprise Monitoring Server and multiple Netcool/OMNIBus ObjectServer

Uses

By configuring the Hub monitoring server to send all events to multiple Netcool/OMNIBus ObjectServers, you can introduce redundancy into your event synchronization environment. Use this method to increase availability by having your event data on multiple Netcool/OMNIBus ObjectServers.

This scenario can also be used to organize event data by functional area. For example, operating system and virtualization situation events can be configured to forward to one Netcool/OMNIBus event server, while WebSphere Application Server, DB2, and other middleware situation events can be configured to forward to another Netcool/OMNIBus event server.

Installation and configuration

If you are installing IBM Tivoli Monitoring and Netcool/OMNIBus event integration for the first time, complete the tasks in this table.

However, if you are upgrading an existing event integration environment, see “Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIbus integration” on page 766.

Table 129. Installation and configuration: one hub Tivoli Enterprise Monitoring Server and multiple Netcool/OMNIbus ObjectServers

	Task	Architecture Type	Administrator
1.	<p>Install Netcool/OMNIbus ObjectServer if it is not already installed or upgrade it to the fix pack version required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIbus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIbus” on page 91.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
2.	<p><i>For each Netcool/OMNIbus ObjectServer:</i> Install the Netcool/OMNIbus Probe for Tivoli EIF if it is not already installed or upgrade it to the release required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIbus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIbus” on page 91.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
3.	<p><i>For each Netcool/OMNIbus ObjectServer:</i> “Installing the IBM Tivoli Monitoring Event Synchronization Component” on page 739.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
4.	<p><i>For each Netcool/OMNIbus ObjectServer:</i> “Updating the OMNIbus database schema on single-tier or aggregation tier ObjectServers” on page 753.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
5.	<p><i>For each Netcool/OMNIbus ObjectServer:</i> “Changing the default deduplication trigger” on page 755.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
6.	<p><i>For each Netcool/OMNIbus ObjectServer:</i> “Configuring the OMNIbus server for program execution from scripts” on page 763.</p>	Bidirectional	Netcool/OMNIbus

Table 129. Installation and configuration: one hub Tivoli Enterprise Monitoring Server and multiple Netcool/OMNIbus ObjectServers (continued)

	Task	Architecture Type	Administrator
7.	For each Netcool/OMNIbus ObjectServer: Start the IBM Tivoli Monitoring Situation Update Forwarder. For more information, see “Starting and stopping the IBM Tivoli Monitoring Situation Update Forwarder” on page 759.	Bidirectional	Netcool/OMNIbus
8.	For each Netcool/OMNIbus ObjectServer: “Configuring the Netcool/OMNIbus EIF probe” on page 760.	Unidirectional and bidirectional	Netcool/OMNIbus
9.	“Configuring the hub monitoring server to forward events” on page 764 to one of the EIF Probes. This probe is the default EIF receiver.	Unidirectional and bidirectional	IBM Tivoli Monitoring
10.	For each additional Netcool/OMNIbus ObjectServer: Use the tacmd createEventDest command to create an event destination for the EIF probe associated with the ObjectServer. See the <i>IBM Tivoli Monitoring Command Reference</i> for a description.	Unidirectional and bidirectional	IBM Tivoli Monitoring
11.	For each situation, configure which event destinations the events for the situation should be sent to.	Unidirectional and bidirectional	IBM Tivoli Monitoring Operator
12.	“Verifying installation and configuration” on page 765.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIbus
13.	Determine if additional configuration tasks should be performed. See “Customizing Event Integration” on page 772.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIbus

Multiple IBM Tivoli Monitoring agents and one Netcool/OMNIbus ObjectServer

In this scenario, IBM Tivoli Monitoring agents are configured to forward events directly to either the Netcool/OMNIbus Probe for Tivoli EIF or the Netcool/OMNIbus Probe for SNMP.

This architecture eliminates the interaction with the hub monitoring server. If your environment uses firewalls, you must ensure any IBM Tivoli Monitoring agents and Netcool/OMNIbus probes can communicate with each other and are not restricted by the firewall. Because events are forwarded directly to Netcool/OMNIbus by the IBM Tivoli Monitoring agents, no bidirectional communication is required. Any actions taken on events through Netcool/OMNIbus are not forwarded back to the IBM Tivoli Monitoring agent.

The Netcool/OMNIbus Probe for Tivoli EIF supports both SSL and non-SSL connections with monitoring agents. If you want to use an SSL connection, you must have version 12.0 or later of the probe.

You can configure the agents to send heartbeat events at regular intervals to Netcool/OMNIbus. You can enable heartbeat automation in the Netcool/OMNIbus ObjectServer. By using heartbeat automation, a "Heartbeat Missing" event is opened if a heartbeat event is not received from an agent in the expected

time frame. By enabling heartbeat automation, the Netcool/OMNIBus operator knows when a monitoring agent is not available. The heartbeat interval is configurable.

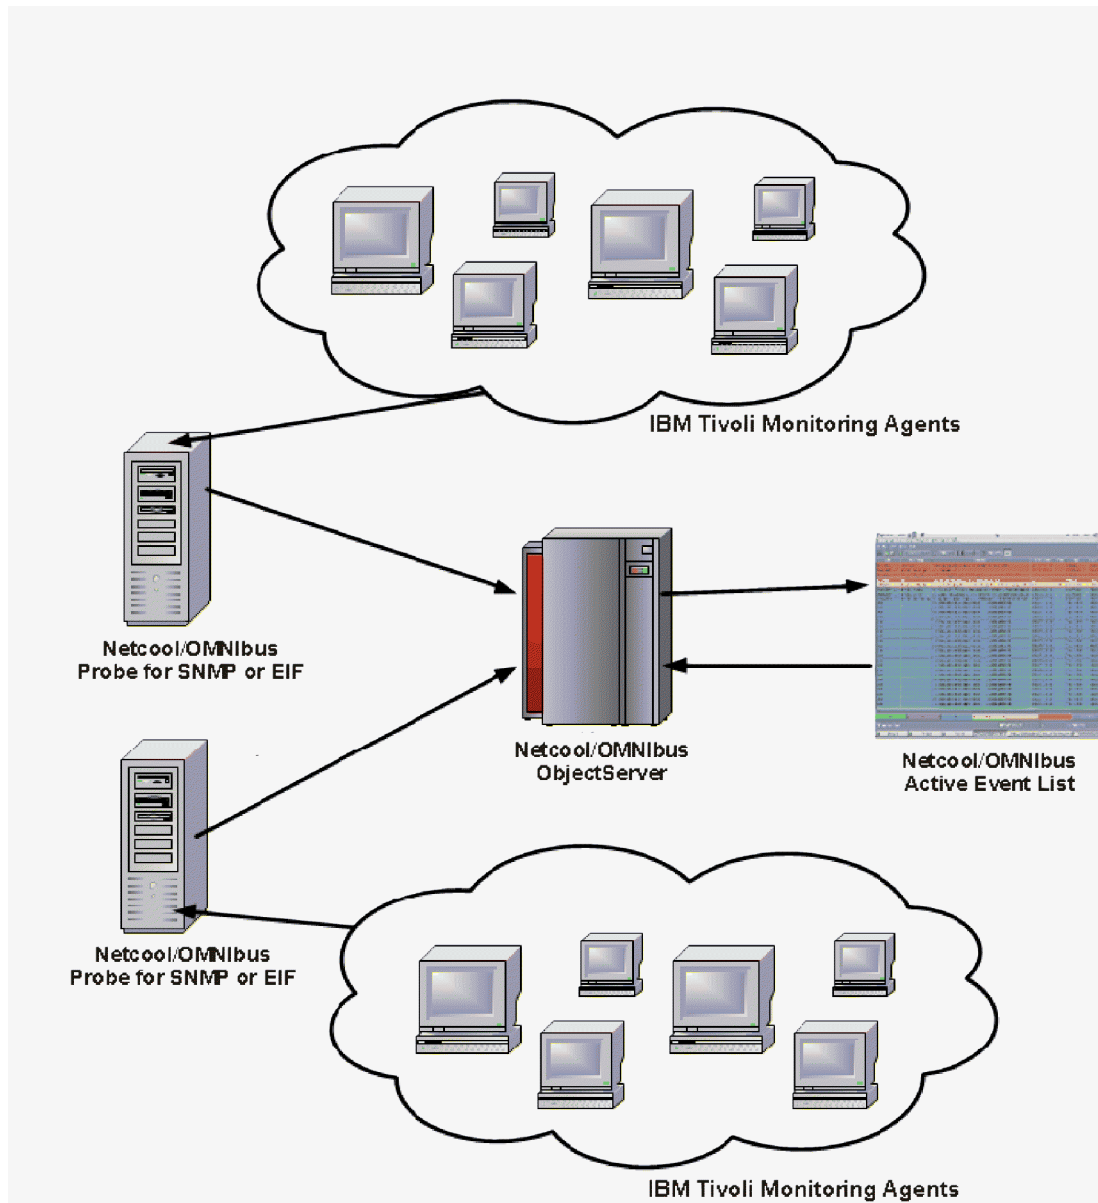


Figure 149. Multiple IBM Tivoli Monitoring Agents and one Netcool/OMNIBus ObjectServer

Uses

Use this scenario when you have agents that are not connected to IBM Tivoli Monitoring Server. You can also use this architecture if you want monitoring server connected agents to send their critical events to the Netcool/OMNIBus Objectserver. The agents can send events directly to the OMNIBus server through the Netcool/OMNIBus Probe for Tivoli EIF or the Probe for SNMP.

Installation and configuration for agents sending events to the Netcool/OMNIBus Probe for Tivoli EIF

If you are installing IBM Tivoli Monitoring and Netcool/OMNIBus event integration for the first time, perform the tasks in this table.

However, if you are upgrading an existing event integration environment, see “Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIbus integration” on page 766.

Note: If agents and a hub monitoring server are sending events to the same Netcool/OMNIbus ObjectServer and probe, you must only perform tasks 1 to 7 once.

Table 130. Installation and configuration for agents sending events to the Netcool/OMNIbus Probe for Tivoli EIF

	Task	Architecture type	Administrator
1.	<p>Install Netcool/OMNIbus ObjectServer if it is not already installed or upgrade it to the fix pack version required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIbus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIbus” on page 91.</p>	Unidirectional	Netcool/OMNIbus
2.	<p>Install the Netcool/OMNIbus Probe for Tivoli EIF if it is not already installed or upgrade it to the release required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIbus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIbus” on page 91.</p>	Unidirectional	Netcool/OMNIbus
3.	“Installing the IBM Tivoli Monitoring Event Synchronization Component” on page 739.	Unidirectional	Netcool/OMNIbus
4.	“Updating the OMNIbus database schema on single-tier or aggregation tier ObjectServers” on page 753.	Unidirectional	Netcool/OMNIbus
5.	“Changing the default deduplication trigger” on page 755.	Unidirectional	Netcool/OMNIbus
6.	Enable heartbeat automation in OMNIbus. For instructions on enabling OMNIbus heartbeat automation, see the <i>Agent Autonomy</i> chapter in the <i>IBM Tivoli Monitoring Administrator's Guide</i> .	Unidirectional	Netcool/OMNIbus

Table 130. Installation and configuration for agents sending events to the Netcool/OMNIbus Probe for Tivoli EIF (continued)

	Task	Architecture type	Administrator
7.	“Configuring the Netcool/OMNIbus EIF probe” on page 760. Note: If an SSL connection will be used between the probe and a monitoring agent, ensure that you have version 12.0 or later of the probe.	Unidirectional	Netcool/OMNIbus
8.	For each IBM Tivoli Monitoring Agent: configure the IBM Tivoli Monitoring Agent to forward events to the EIF Probe. See the <i>Agent autonomy</i> chapter in the <i>IBM Tivoli Monitoring Administrator's Guide</i> for information on configuring agents to send EIF events to OMNIbus using either SSL or non-SSL connections.	Unidirectional	IBM Tivoli Monitoring
9.	“Verifying installation and configuration” on page 765.	Unidirectional	IBM Tivoli Monitoring and Netcool/OMNIbus

Installation and configuration if agents are sending events to the Netcool/OMNIbus Probe for SNMP

If you are installing IBM Tivoli Monitoring and Netcool/OMNIbus event integration for the first time, perform the tasks in this table.

Table 131. Installation and configuration if agents are sending events to the Netcool/OMNIbus Probe for SNMP

	Task	Architecture Type	Administrator
1.	Install Netcool/OMNIbus ObjectServer if it is not already installed or upgrade it to the fix pack version required by IBM Tivoli Monitoring. See the Netcool/OMNIbus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm . See also “Required software for event integration with Netcool/OMNIbus” on page 91.	Unidirectional	Netcool/OMNIbus

Table 131. Installation and configuration if agents are sending events to the Netcool/OMNIBus Probe for SNMP (continued)

	Task	Architecture Type	Administrator
2.	<p>Install the Netcool/OMNIBus Probe for SNMP if it is not already installed or upgrade it to the release required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIBus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIBus” on page 91.</p>	Unidirectional	Netcool/OMNIBus
3.	<p>Enable heartbeat automation in OMNIBus. For instructions on enabling OMNIBus heartbeat automation, see the <i>Agent Autonomy</i> chapter in the <i>IBM Tivoli Monitoring Administrator's Guide</i>.</p>	Unidirectional	Netcool/OMNIBus
4.	<p>Configure the Netcool/OMNIBus Probe for SNMP to process IBM Tivoli Monitoring agent events. For instructions on configuring OMNIBus to receive SNMP alerts, see the <i>Agent Autonomy</i> chapter in the <i>IBM Tivoli Monitoring Administrator's Guide</i>.</p>	Unidirectional	Netcool/OMNIBus
5.	<p>For each IBM Tivoli Monitoring Agent: configure the IBM Tivoli Monitoring Agent to forward events to the SNMP Probe. See the <i>Agent Autonomy</i> chapter in the <i>IBM Tivoli Monitoring Administrator's Guide</i> for information on configuring agents to send events to OMNIBus.</p>	Unidirectional	IBM Tivoli Monitoring
6.	<p>“Verifying installation and configuration” on page 765.</p>	Unidirectional	IBM Tivoli Monitoring and Netcool/OMNIBus

Netcool/OMNIBus Multitiered and High-availability Architecture

The IBM Tivoli Netcool/OMNIBus product can be deployed in a multitiered configuration to increase performance and event handling capacity.

High-availability is also supported by adding primary and backup servers to the tiers. In the standard multitiered environment, three sets of Netcool/OMNIBus ObjectServers are included:

1. The collection tier includes the ObjectServers to which probes connect. In a high-availability architecture, primary and backup pairs of ObjectServers are included in this tier. Each ObjectServer in the collection tier is connected to an ObjectServer in the aggregation tier by using a unidirectional gateway.

2. The aggregation tier includes up to two ObjectServers if you are using a high-availability architecture. The ObjectServers are connected by a bidirectional gateway. The bulk of the event processing occurs in the aggregation tier.
3. The display tier includes one or more ObjectServers. They are connected to an ObjectServer in the aggregation tier by using a unidirectional gateway. The Netcool/OMNIBus desktop event list users and web GUI users connect to ObjectServers in this tier. The operator connects their Netcool/OMNIBus Desktop UI to the display tier in dual-server desktop mode. Events are retrieved from the display ObjectServers, but updates made to events go to both the display and the aggregation tiers.

High-availability can also be achieved in a single-tier environment by having a primary and backup ObjectServer and a bidirectional gateway between them. In a single-tier architecture, the probes and user interfaces connect directly to these ObjectServers and the ObjectServers perform the collection, aggregation, and display functions.

The IBM Tivoli Monitoring triggers are assigned to the `primary_only` trigger group. In a high-availability Netcool/OMNIBus architecture, the standard Netcool/OMNIBus automations enable the triggers in the `primary_only` trigger group on the acting primary ObjectServer and disable them on the backup ObjectServer.

For more information about the Netcool/OMNIBus ObjectServer multitiered and high-availability architecture and setup instructions, see the Netcool/OMNIBus information center: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm.

You can also set up peer-to-peer failover mode for Netcool/OMNIBus probes to reduce event loss. Two instances of a probe can run simultaneously in a peer-to-peer failover relationship. One instance is designated as the master; the other instance acts as a slave and is on hot standby. If the master instance fails, the slave instance is activated. For more information about configuring and using master and slave probes for high availability, see the Netcool/OMNIBus information center: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm.

In Netcool/OMNIBus multitier and high-availability architectures, the Tivoli Enterprise Monitoring Servers are configured to forward events to the ObjectServers via the probes. If probe failover mode is used, you must configure situation events in IBM Tivoli Monitoring with multiple event destinations – one destination to the master probe and a second destination to the slave probe.

In a multitiered architecture, IBM Tivoli Monitoring provides triggers to add to the collection tier, database schema updates and triggers for the aggregation tier, and database schema updates for the display tier. The gateways also must be updated with mapping entries. In this architecture, the IBM Tivoli Monitoring Situation Update Forwarder is installed on each ObjectServer in the aggregation tier.

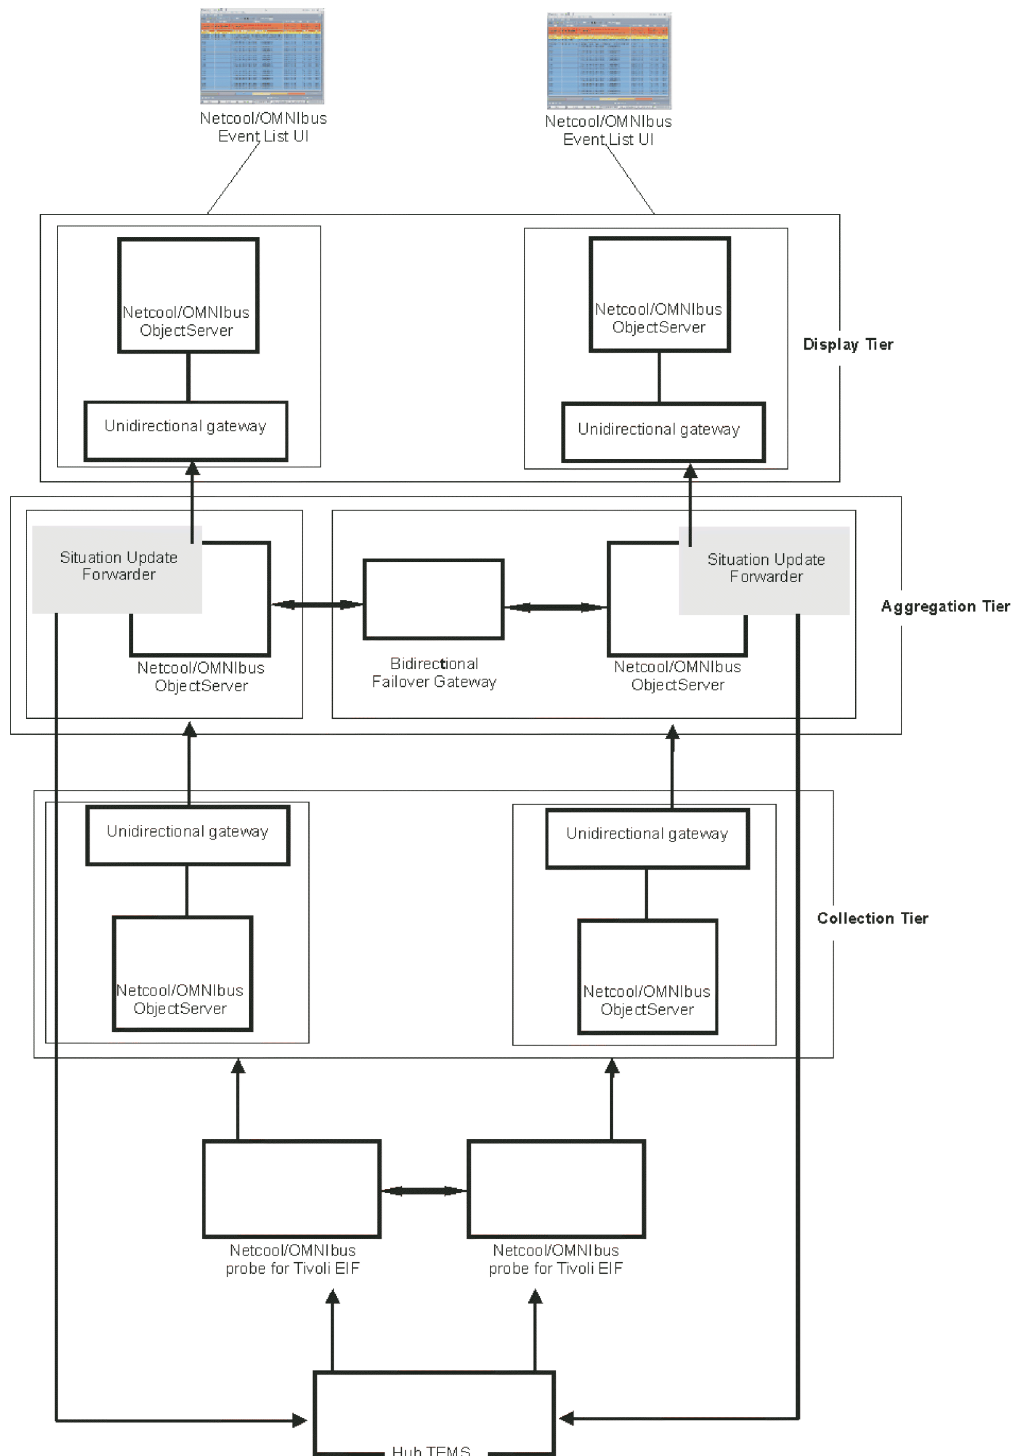


Figure 150. Standard multitiered architecture with high availability

In a single-tier high-availability architecture, the IBM Tivoli monitoring triggers and database schema updates are applied to the primary and backup ObjectServers and mapping entries are configured for the bidirectional failover gateway between the ObjectServers. The IBM Tivoli Monitoring Situation Update Forwarder is installed on each ObjectServer.

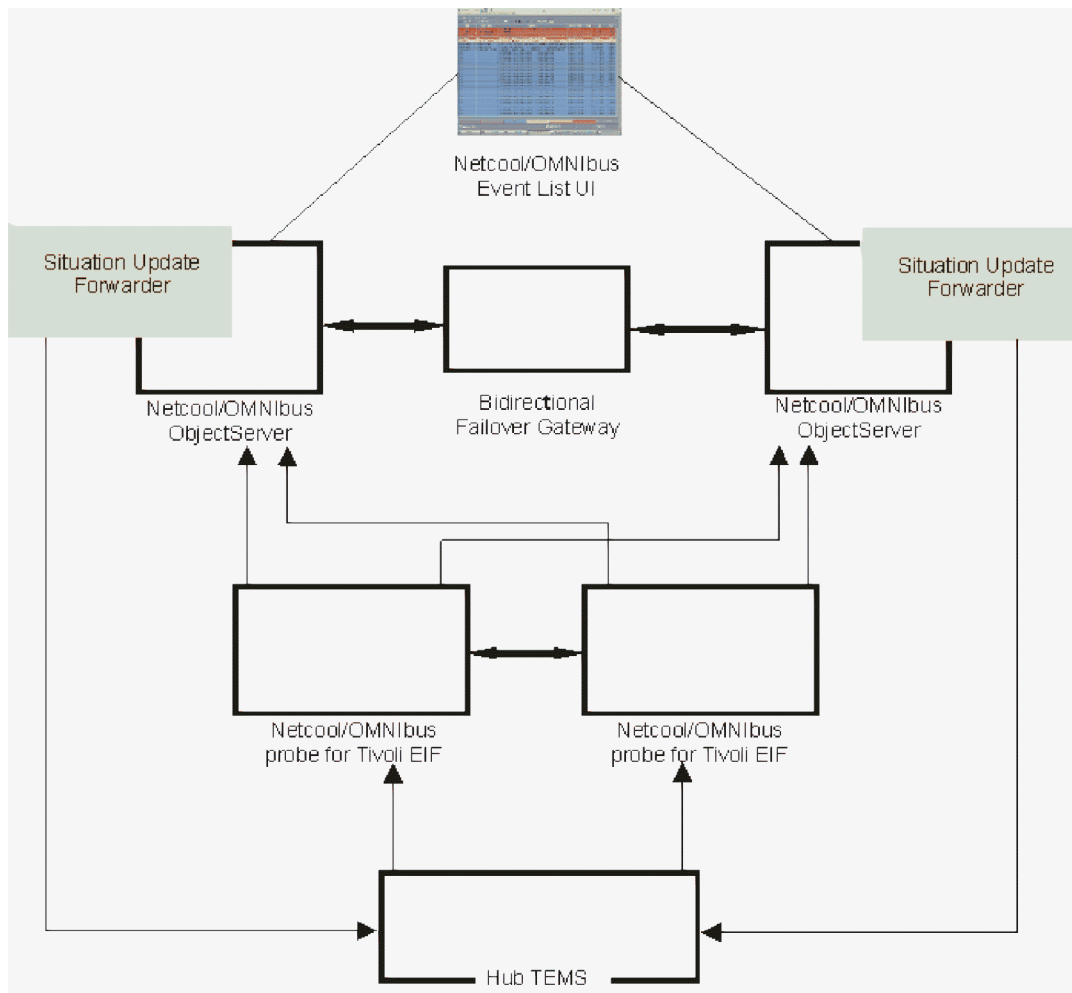


Figure 151. single-tier high-availability architecture

Uses

You can use the Netcool/OMNIBus multitiered architecture if you need to add scalability to your Netcool/OMNIBus environment to increase performance and event handling. High availability can be added for the ObjectServers in your single-tier and multitier architectures so that your environment can continue to operate at full capacity (and with minimal event loss) in the event of ObjectServer failure.

Adding peer-to-peer failover for the probes can be used to reduce event loss in the event of a probe failure.

Installation and configuration

If you are installing IBM Tivoli Monitoring and Netcool/OMNIBus event integration for the first time, perform the tasks in this table.

However, if you are upgrading an existing event integration environment, see “Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIBus integration” on page 766.

Table 132. Installation and configuration: Netcool/OMNIbus Multitiered and High-Availability Architecture ObjectServers

	Task	Architecture Type	Administrator
1.	<p>Install Netcool/OMNIbus ObjectServers in a multitiered architecture or single-tier high availability environment if they are not already installed, or upgrade them to the fix pack version required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIbus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIbus” on page 91.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
2.	<p>Install the Netcool/OMNIbus Probe for Tivoli EIF if it is not already installed or upgrade it to the release required by IBM Tivoli Monitoring. If you are using peer-to-peer probe failover, install and configure master and slave EIF probes.</p> <p>See the Netcool/OMNIbus Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIbus” on page 91.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
3.	<p>For each Netcool/OMNIbus ObjectServer in the aggregation tier or for each ObjectServer in a single-tier architecture, see “Installing the IBM Tivoli Monitoring Event Synchronization Component” on page 739.</p> <p>Note: If you are installing Netcool/OMNIbus ObjectServer in a high availability environment, you must install the event synchronization component on the primary and backup ObjectServer in the single-tier architecture and in the aggregation tier.</p>	Unidirectional and bidirectional	Netcool/OMNIbus

Table 132. Installation and configuration: Netcool/OMNIbus Multitiered and High-Availability Architecture ObjectServers (continued)

	Task	Architecture Type	Administrator
4.	For each gateway between ObjectServers in a multitier or single-tier architecture, see “Updating gateways to map attributes” on page 750.	Unidirectional and bidirectional	Netcool/OMNIbus
5.	For each bidirectional failover gateway in an aggregation tier or single-tier with high availability, see “Updating the bidirectional failover gateway to replicate tables” on page 751.	Unidirectional and bidirectional	Netcool/OMNIbus
6.	For each Netcool/OMNIbus ObjectServer in the collection tier: “Updating the OMNIbus database schema in the collection tier” on page 752.	Unidirectional and bidirectional	Netcool/OMNIbus
7.	For each Netcool/OMNIbus ObjectServer in the aggregation tier or for each ObjectServer in the single-tier architecture, see “Updating the OMNIbus database schema on single-tier or aggregation tier ObjectServers” on page 753.	Unidirectional and bidirectional	Netcool/OMNIbus
8.	For each ObjectServer in the aggregation tier or in a single-tier architecture, see “Changing the default deduplication trigger” on page 755	Unidirectional and bidirectional	Netcool/OMNIbus
9.	For each Netcool/OMNIbus ObjectServer in the aggregation tier or in the single-tier architecture, see “Configuring the OMNIbus server for program execution from scripts” on page 763.	Bidirectional	Netcool/OMNIbus
10.	For each Netcool/OMNIbus ObjectServer in the aggregation tier or in the single-tier architecture, Start the IBM Tivoli Monitoring Situation Update Forwarder. For more information, see “Starting and stopping the IBM Tivoli Monitoring Situation Update Forwarder” on page 759.	Bidirectional	Netcool/OMNIbus
11.	For each Netcool/OMNIbus ObjectServer in the display tier, see “Updating the OMNIbus database schema in the display tier” on page 758.	Unidirectional and bidirectional	Netcool/OMNIbus

Table 132. Installation and configuration: Netcool/OMNIBus Multitiered and High-Availability Architecture ObjectServers (continued)

	Task	Architecture Type	Administrator
12.	Restart each gateway that was updated in steps 4 and 5 based on the following suggested restart order for a multitier architecture: 1. Restart gateways between the collection and aggregation tiers. 2. Restart bidirectional failover gateways in the aggregation tier. 3. Restart gateways between the aggregation and display tiers. See the Netcool/OMNIBus Information Center for more details on restarting the gateways: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm .	Unidirectional and bidirectional	Netcool/OMNIBus
13.	For each EIF probe, see “Configuring the Netcool/OMNIBus EIF probe” on page 760. Note: If you are using a multitier architecture you must also ensure that the EIF probes are configured to connect to an ObjectServer in the collection tier.	Unidirectional and bidirectional	Netcool/OMNIBus
14.	For each hub monitoring server, see “Configuring the hub monitoring server to forward events” on page 764. If you are using peer-to-peer probe failover, configure the Hub to forward events to the master probe.	Unidirectional and bidirectional	IBM Tivoli Monitoring
15.	If you are using peer-to-peer probe failover, use the tacmd createEventDest command to create an event destination for the EIF slave probe. For more information, see the <i>IBM Tivoli Monitoring Command Reference</i> .	Unidirectional and bidirectional	IBM Tivoli Monitoring
16.	If you are using peer-to-peer probe failover, configure each situation to be forwarded to the event destination for the master probe and to the event destination for the slave probe using the Tivoli Enterprise Portal Situation Editor or the tacmd createsit or editsit commands.	Unidirectional and bidirectional	IBM Tivoli Monitoring
17.	“Verifying installation and configuration” on page 765.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIBus

Table 132. Installation and configuration: Netcool/OMNIBus Multitiered and High-Availability Architecture ObjectServers (continued)

	Task	Architecture Type	Administrator
18.	Determine if additional configuration tasks should be performed. See “Customizing Event Integration” on page 772.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIBus

Integration with Tivoli Business Service Manager

In this scenario, IBM Tivoli Monitoring is configured to forward events to Netcool/OMNIBus. The Tivoli Business Service Manager Data Server analyzes the events in the Netcool/OMNIBus ObjectServer for matches against the incoming-status rules configured for service models. If the event's data matches a status rule, the status of the service model is changed accordingly.

Tivoli Business Service Manager can be used to install Netcool/OMNIBus ObjectServer and the Netcool/OMNIBus Probe for Tivoli EIF if you do not have an existing installation of these components. See the Tivoli Business Service Manager information center for details.

You can use any of the architecture scenarios described in this chapter for configuring forwarding of events from IBM Tivoli Monitoring to Netcool/OMNIBus when integrating with Tivoli Business Service Manager.




Figure 152. Integrating IBM Tivoli Monitoring, Netcool/OMNIBus, and Tivoli Business Service Manager

Uses

If you are using Tivoli Business Service Manager to manage your business services, and you want to use monitoring events to update business service status, then you must use the IBM Tivoli Monitoring probe rules that set the OMNIBus attributes required by Tivoli Business Service Manager.

Installation and configuration

If you are installing IBM Tivoli Monitoring and Netcool/OMNIBus event integration for the first time, perform the tasks in this table.

However, if you are upgrading an existing event integration environment, see “Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIBus integration” on page 766.

You should use the IBM Tivoli Monitoring probe rules file and .sql files provided with IBM Tivoli Monitoring instead of the related Tivoli Monitoring rules file and .sql files provided with Tivoli Business Service Manager Version 4.2.1.

Table 133. Installation and configuration

	Task	Architecture Type	Administrator
1.	<p>Install Netcool/OMNIbus ObjectServer using the Tivoli Business Service Manager installation program or follow the procedures in the Tivoli Business Service Manager documentation for updating an existing ObjectServer with database schema updates required by Tivoli Business Service Manager.</p> <p>See the Tivoli Business Service Manager Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/index.jsp?topic=/com.ibm.tivoli.itbsm.doc/welcome.htm.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
2.	<p>Install the Netcool/OMNIbus Probe for Tivoli EIF using the Tivoli Business Service Manager installation program or follow the procedures in the Tivoli Business Service Manager documentation for updating an existing EIF probe.</p> <p>See the Tivoli Business Service Manager Information Center for detailed instructions on this task: http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/index.jsp?topic=/com.ibm.tivoli.itbsm.doc/welcome.htm.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
3.	"Installing the IBM Tivoli Monitoring Event Synchronization Component" on page 739.	Unidirectional and bidirectional	Netcool/OMNIbus
4.	"Updating the OMNIbus database schema on single-tier or aggregation tier ObjectServers" on page 753.	Unidirectional and bidirectional	Netcool/OMNIbus
5.	"Changing the default deduplication trigger" on page 755.	Unidirectional and bidirectional	Netcool/OMNIbus
6.	"Configuring the OMNIbus server for program execution from scripts" on page 763.	Bidirectional	Netcool/OMNIbus
7.	Start the IBM Tivoli Monitoring Situation Update Forwarder. For more information, see "Starting and stopping the IBM Tivoli Monitoring Situation Update Forwarder" on page 759.	Bidirectional	Netcool/OMNIbus

Table 133. Installation and configuration (continued)

	Task	Architecture Type	Administrator
8.	If you have more than one monitoring server forwarding events to the Netcool/OMNIbus ObjectServer, define additional monitoring servers to the IBM Tivoli Monitoring Situation Update Forwarder. For more information, see “Updating the IBM Tivoli Monitoring Situation Forwarder to forward event status updates to additional monitoring servers” on page 774.	Bidirectional	Netcool/OMNIbus
9.	To use the IBM Tivoli Monitoring probe rules include files for Tivoli Business Service Manager integration, see “Configuring the Netcool/OMNIbus EIF probe” on page 760. Note: Some monitoring agents also have their own probe rules include file. See the user guide for each monitoring agent that you plan to install to determine if it has a rules include file that should be copied to the computer system where the probe is installed and uncommented in the probe's main rules file, <code>tivoli_eif.rules</code> .	Unidirectional and bidirectional	Netcool/OMNIbus
10.	For each hub monitoring server: “Configuring the hub monitoring server to forward events” on page 764.	Unidirectional and bidirectional	IBM Tivoli Monitoring
11.	“Verifying installation and configuration” on page 765.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIbus
12.	Determine if additional configuration tasks should be performed. See “Customizing Event Integration” on page 772.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIbus

Installation and configuration

The products in this section must be installed and configured before you install the event synchronization component and configure event forwarding for Netcool/OMNIbus.

- IBM Tivoli Netcool/OMNIbus V7.2.1 or later
- IBM Tivoli Netcool/OMNIbus probe for Tivoli EIF version 10 or later and the non-native probe version 12 or later
- IBM Tivoli Monitoring

Note: fix packs might be required for each of the supported Netcool/OMNIbus releases, for more information see “Required software for event integration with Netcool/OMNIbus” on page 91.

The steps for installing and configuring event integration depend on the type of architecture you are installing. The section on “Architecture scenarios” on page 715 describes various architecture types and steps to implement the solution. First select the type of architecture you want to implement, and then follow the steps outlined for that architecture type.

Setting up event forwarding between IBM Tivoli Monitoring and Netcool/OMNIBus involves both IBM Tivoli Monitoring and Netcool/OMNIBus configuration tasks, therefore setup should be coordinated between both the IBM Tivoli Monitoring and Netcool/OMNIBus administrators.

Installing the IBM Tivoli Monitoring Event Synchronization Component

Installing the IBM Tivoli Monitoring Event Synchronization component also installs the IBM Tivoli Monitoring Situation Update Forwarder, IBM Tivoli Monitoring rules file for the Netcool/OMNIBus Probe for Tivoli EIF, and SQL files to update the Netcool/OMNIBus ObjectServer database for EIF event handling. On Windows, a Situation Update Forwarder service is also created. The tasks in this section should be completed by the Netcool/OMNIBus administrator.

The IBM Tivoli Monitoring Event Synchronization Component installer is located on the *IBM Tivoli Monitoring Tools* DVD or DVD image.

- **Bidirectional architecture:**

Install the IBM Tivoli Monitoring Event Synchronization Component on the host of the Netcool/OMNIBus ObjectServer. The bidirectional architecture requires the IBM Tivoli Monitoring Situation Update Forwarder to be located on the Netcool/OMNIBus ObjectServer.

- **Unidirectional architecture:**

A unidirectional architecture does not need the IBM Tivoli Monitoring Situation Update Forwarder, therefore the IBM Tivoli Monitoring Event Synchronization Component can be installed on any system. The files required for later setup steps, such as probe rules and database SQL files, can be copied manually to the Netcool/OMNIBus Probe for Tivoli EIF and Netcool/OMNIBus ObjectServer. During installation you are asked questions that only apply to the bidirectional architecture. You must specify a valid value for these parameters even though they will not be used.

If you are using a Netcool/OMNIBus multitier architecture, install the Tivoli Monitoring Event Synchronization component with each ObjectServer in the aggregation tier. If you are using a single-tier architecture, install the event synchronization component with each of your ObjectServers.

Use one of the following three methods to install the Tivoli Monitoring Event Synchronization component:

- Installing from a wizard
- Installing from the command-line
- Installing from the command-line using a silent installation

Note:

- See “Required software for event integration with Netcool/OMNIBus” on page 91 to determine if there are any operating system library requirements for the system on which you are installing the event synchronization component.
- See the *Event synchronization installation and configuration troubleshooting* section in the *IBM Tivoli Monitoring Troubleshooting Guide* if you encounter any errors during the installation process.
- You cannot install event synchronization for Netcool/OMNIBus on the same system as an IBM Tivoli Enterprise Console event server.
- If the ObjectServer is running on Windows 2003 and you are planning to install the event synchronization remotely (using a program such as Terminal Services to connect to that Windows 2003 computer), you must run the **change user /install** command before you run the installation, which puts the computer into the required “install” mode. After the installation, run the **change user /execute** command to return the computer to its previous mode.

- If you have a monitoring server on an operating system like UNIX or Linux, you must configure your TCP/IP network services in the /etc/hosts file to return the fully qualified host name if your Netcool/OMNIBus ObjectServer must use the fully qualified hostname to send event status updates to the monitoring server.
- Linux or UNIX users can run the event synchronization installer under a root or a non-root user ID. AIX users must run the event synchronization installer under a root user ID.
- The results of the event synchronization installation are written to the %TEMP%\itm_tec_event_sync_install.log file on Windows and to the /tmp/itm_tec_event_sync_install.log file on UNIX.

Installing from a wizard

This section outlines the steps for installing event synchronization from a wizard.

Procedure

1. On the computer where the ObjectServer is installed, launch the event synchronization installation:
 - **Windows** Double-click the ESync3000Win32.exe file in the tec subdirectory on the IBM Tivoli Monitoring Tools DVD or DVD image.
 - **Linux** **UNIX** Change to the tec subdirectory on the IBM Tivoli Monitoring Tools DVD or DVD image and run the following command:
ESync3000operating_system.bin

where *operating_system* is the operating system you are installing on (Aix, HP11, Linux, linux390, or Solaris). For example, run the following command on an AIX computer:

ESync3000Aix.bin

If the installer cannot locate OMNIBus in its usual place, the following window is displayed. Click **Next** to continue installing the event synchronization:

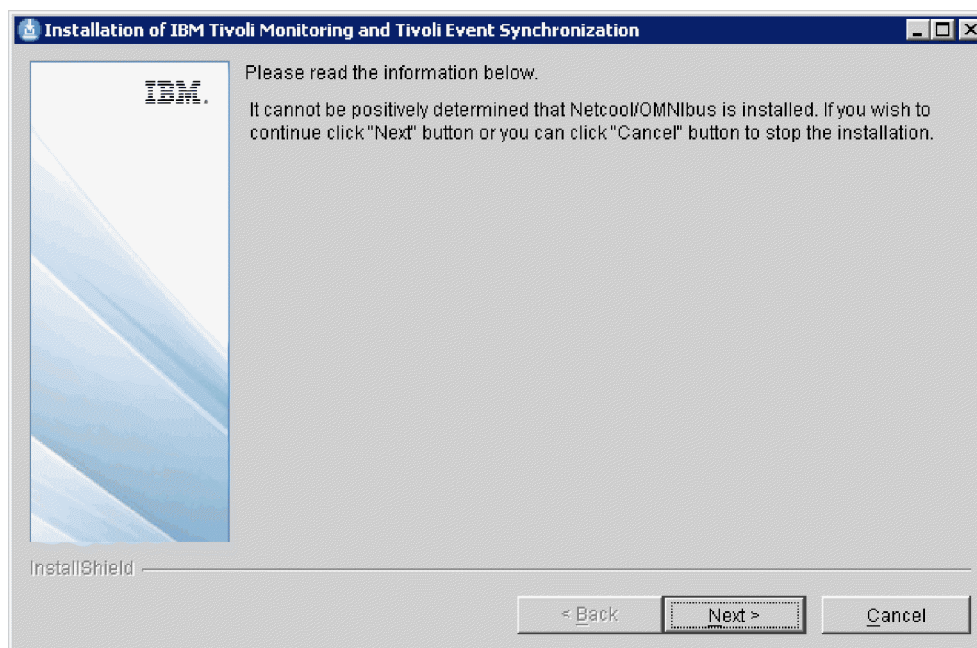


Figure 153. Installation of IBM Tivoli Monitoring and Tivoli Event Synchronization

2. Click **Next** on the Welcome window.
3. Review the license agreement, select **I accept the terms in the license agreement** and click **Next**.

4. Click **Next** to install the synchronization component in the default location, or use the **Browse** button to select another location.
 - **Windows** The default installation directory is C:\Program Files\IBM\SitForwarder.
 - **Linux** **UNIX** The default installation directory is /opt/IBM/SitForwarder.

Click **Next** to continue.

5. Complete the fields in the installation windows by using the configuration values described in Table 134 and click **Next**.

The screenshot shows a Windows-style installation window titled "Installation of IBM Tivoli Monitoring and Tivoli Event Synchronization". On the left is the IBM logo. The main area contains several configuration fields with labels and input boxes or dropdowns:

- Name of configuration file:** situpdate.conf
- Number of seconds to sleep when no new situation updates:** 3
- Number of bytes to use to save last event:** 50
- URL of the TEMS SOAP server:** cms/soap
- Rate for sending SOAP requests to TEMS from Event Sync via Web Services:** 10
- Level of debug detail for log:** low (selected from a dropdown menu)

At the bottom, there are three buttons: "< Back", "Next >" (which is highlighted with a dashed border), and "Cancel". The "InstallShield" logo is visible in the bottom left corner.

Figure 154. Netcool/OMNIBus event synchronization configuration fields

Table 134. Netcool/OMNIBus event synchronization configuration fields

Field	Description
Name of configuration file	<i>For bidirectional architecture only:</i> The name of the file where the IBM Tivoli Monitoring Situation Update Forwarder configuration information is stored. The default name is situpdate.conf.
Number of seconds to sleep when no new situation updates	<i>For bidirectional architecture only:</i> The polling interval, in seconds used by the IBM Tivoli Monitoring Situation Update Forwarder to determine if there are new event updates to send from Netcool/OMNIBus to the hub Tivoli Enterprise Monitoring Server. The minimum value is 1, while the default value is 3. If no situation events are found, the IBM Tivoli Monitoring Situation Update Forwarder rests for 3 seconds.
Number of bytes to use to save last event	<i>For bidirectional architecture only:</i> Number of bytes that the IBM Tivoli Monitoring Situation Update Forwarder uses when it saves the location of the last event it processes. This value must be an integer. The minimum (and default) value is 50.

Table 134. Netcool/OMNIBus event synchronization configuration fields (continued)

Field	Description
URL of the Tivoli Enterprise Monitoring Server SOAP Server	<i>For bidirectional architecture only:</i> The URL for the SOAP Server configured on the computer where the hub monitoring server is running. The default value is cms/soap. Do not change this value. The actual URL for the SOAP Server is dynamically constructed from the incoming event.
Rate for sending SOAP requests to the monitoring server from the OMNIBus probe via web services	<i>For bidirectional architecture only:</i> The maximum number of event updates sent by the IBM Tivoli Monitoring Situation Update Forwarder to the hub monitoring server at one time. The minimum (and default) value is 10 events.
Level of debug detail for log	<i>For bidirectional architecture only:</i> The level of information for event synchronization that is logged. The following choices are available: <ul style="list-style-type: none"> • Low (default) • Medium • Verbose

6. Complete the fields on the installation window by using the values described in Table 135 and click **Next**.

Table 135. Netcool/OMNIBus event synchronization configuration fields, continued

Field	Description
Maximum size of any single cache file	<i>For bidirectional architecture only:</i> The maximum permitted size, in bytes, for any one event cache file used by the IBM Tivoli Monitoring Situation Update Forwarder. The minimum (and default) value is 50000. Do not use commas when specifying this value (specify 50000 instead of 50,000).
Maximum number of caches files	<i>For bidirectional architecture only:</i> The maximum number of event caches files at any given time used by the IBM Tivoli Monitoring Situation Update Forwarder. The minimum value is 2, while the default value is 10. When this value is reached, the oldest file is deleted to make room for a new file.
Directory for cache files to be located	<i>For bidirectional architecture only:</i> The location where IBM Tivoli Monitoring Situation Update Forwarder event cache files are located. The default locations are as follows: <ul style="list-style-type: none"> • Windows C:\Program Files\IBM\SitForwarder\persistence. • UNIX /opt/IBM/TEC/SitForwarder/persistence

7. Type the following information for each hub monitoring server for which you want to use bidirectional synchronization and click **Add**. You must specify information for at least one hub monitoring server even if you are only using unidirectional synchronization.

Host name

The fully qualified host name for the computer where the hub monitoring server is running. The name must match the information that will be in events coming from this hub monitoring server.

User ID

The user ID to access the computer where the hub monitoring server is running.

Password

The password to access the computer.

Confirmation

The password, again.

Repeat this step to add a short host name for the same hub monitoring server by specifying the short host name value for the Host name parameter. You can add information for up to 10 hub monitoring servers host names in this wizard. If you want to add additional monitoring servers, do so after installation by using the steps outlined in “Defining additional monitoring servers to the event server” on page 822. If you have configured the Hot Standby feature, you must configure the host name information for both the primary and secondary hub monitoring servers.

8. When you have provided information about all of the monitoring servers, click **Next**.
A summary window is displayed.
9. Click **Next** to proceed. The installation begins and a progress indicator is displayed. When the installation is completed, a message is displayed telling you that the installation has been successful.
10. Click **Finish** to exit the wizard.

What to do next

If any configuration errors occurred during installation and configuration, you are directed to a log file that contains additional troubleshooting information.

Installing from the command-line

This section outlines the steps for installing event synchronization from the command-line.

Procedure

1. Change to the `tec` directory on the IBM Tivoli Monitoring Tools DVD or DVD image.
2. Run the following command to launch the installation:

```
Windows
ESync3000Win32.exe -console

Linux  UNIX
ESync3000operating_system.bin -console
```

where *operating_system* is the operating system you are installing on (Aix, HP11, Linux, linux390, or Solaris). For example, run the following command on an AIX computer:

```
ESync3000Aix.bin -console
```

The following prompt is displayed:

```
Press 1 for Next, 3 to Cancel or 4 to Redisplay [1]
```

3. Type 1 to start the installation and press Enter.

The following prompt is displayed:

Software Licensing Agreement:

Press Enter to display the license agreement on your screen. Please read the agreement carefully before installing the Program. After reading the agreement, you will be given the opportunity to accept it or decline it. If you choose to decline the agreement, installation will not be completed and you will not be able to use the Program.

4. Press Enter to display the software license agreement.
5. Type 1 and press Enter to accept the license.

The following prompt is displayed:

```
Press 1 for Next, 2 for Previous, 3 to Cancel, or 4 to Redisplay [1]
```

6. Type 1 and press Enter to continue.

The following prompt is displayed:

Name of configuration file [situpdate.conf]

7. Press Enter to use the default configuration file situpdate.conf for the IBM Tivoli Monitoring Situation Update Forwarder. If you want to use a different configuration file, type the name and press Enter.

The following prompt is displayed:

Number of seconds to sleep when no new situation updates [3]

8. Type the number of seconds that you want to use for the polling interval used by the IBM Tivoli Monitoring Situation Update Forwarder to determine if there are new event updates to send from Netcool/OMNIBus to the hub Tivoli Enterprise Monitoring Server and press Enter. The default value is 3, while the minimum value is 1. This configuration value is used for bidirectional architecture only.

The following prompt is displayed:

Number of bytes to use to save last event [50]

9. Type the number of bytes the IBM Tivoli Monitoring Situation Update Forwarder uses to save the last event and press Enter. The default and minimum value is 50. This configuration value is used for bidirectional architecture only.

The following prompt is displayed:

URL of the CMS SOAP server [cms/soap]

10. Type the URL for the hub monitoring server SOAP Server and press Enter. The default value is cms/soap (which you can use if you set up your monitoring server using the defaults for SOAP server configuration). This configuration value is used for bidirectional architecture only.

The following prompt is displayed:

Rate for sending SOAP requests to CMS from TEC via Web Services [10]

11. Type maximum number of event updates to be sent by the IBM Tivoli Monitoring Situation Update Forwarder to the hub monitoring server at one time and press Enter. The default and minimum value is 10. This configuration value is used for bidirectional architecture only.

The following prompt is displayed:

Level of debug for log

```
[x] 1 low
[ ] 2 med
[ ] 3 verbose
```

To select an item enter its number, or enter 0 when you are finished: [0]

12. Type the level of debugging that you want to use and press Enter. The default is Low, indicated by an x next to Low.
13. Type 0 when you have finished and press Enter.

The following prompt is displayed:

Press 1 for Next, 2 for Previous, 3 to Cancel, or 4 to Redisplay [1]

14. Type 1 and press Enter to continue.

The following prompt is displayed:

Maximum size of any single cache file, in bytes [50000]

15. Type the maximum size, in bytes, for the cache file used by the IBM Tivoli Monitoring Situation Update Forwarder and press Enter. The default is 50000. Do not use commas (,) when specifying this value. This configuration value is used for bidirectional architecture only.

The following prompt is displayed:

Maximum number of cache files [10]

16. Type the maximum number of cache files used by the IBM Tivoli Monitoring Situation Update Forwarder to have at one time and press Enter. The default is 10, while the minimum is 2. This configuration value is used for bidirectional architecture only.

 The following prompt is displayed:

Directory for cache files to reside [C:/Program Files/IBM/SitForwarder/persistence]

UNIX The following prompt is displayed:

Directory for cache files to reside [/opt/IBM/SitForwarder/persistence]

17. Type the directory for the IBM Tivoli Monitoring Situation Update Forwarder cache files and press Enter. The default directory is

Windows C:\Program Files\IBM\SitForwarder\persistence

UNIX /opt/IBM/SitForwarder/persistence

The following prompt is displayed:

Press 1 for Next, 2 for Previous, 3 to Cancel, or 4 to Redisplay [1]

This configuration value is used for bidirectional architecture only.

18. Type 1 and press Enter to continue.

19. The following prompt is displayed:

--- Tivoli Enterprise Monitoring Server 1 ---

Host name []

Type the fully qualified host name for the computer where the hub monitoring server is running. This should match the information that will be in events coming from this monitoring server. This configuration value is used for bidirectional architecture only. Press Enter.

The following prompt is displayed:

User ID []

20. Type the user ID to use to access the computer where the hub monitoring server is running. This configuration value is used for bidirectional architecture only. Press Enter.

The following prompt is displayed:

Password:

21. Type the password to access the hub monitoring server. This configuration value is used for bidirectional architecture only. Press Enter.

The following prompt is displayed:

Confirmation:

22. Type the password again to confirm it and press Enter.

The following prompt is displayed:

--- Tivoli Enterprise Monitoring Server 2 ---

Host name []

23. Repeat steps 19 - 22 to add the short host name for the same hub monitoring server. Then repeat the same steps for each additional hub monitoring server that the IBM Tivoli Monitoring Situation Update Forwarder should send event status updates to. If you have configured the Hot Standby feature, you must configure the host name information for both the primary and secondary hub monitoring servers.

When you have provided information for all of the hub monitoring servers *and* you specified information for less than 10 monitoring server host names, press Enter to move through the remaining fields defining additional monitoring servers. Do not specify any additional monitoring server information.

24. When you see the following prompt, type 1 and press Enter to continue:

Press 1 for Next, 2 for Previous, 3 to cancel or 4 to Redisplay [1]

The event synchronization is installed. The following prompt is displayed:

IBM Tivoli Monitoring and Tivoli Enterprise Console Event
Synchronization will
be installed in the following location:

```
/opt/IBM/SitForwarder
```

```
for a total size:
```

```
101.3 MB
```

```
Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]
```

25. Type 1 and press Enter to continue.

The following prompt is displayed:

```
The InstallShield Wizard has successfully installed IBM Tivoli  
Monitoring and Tivoli Event Synchronization.
```

```
Choose Finish to exit the wizard.
```

```
Press 3 to Finish, or 4 to Redisplay [1]
```

26. Type 3 to finish and press Enter.

Installing from the command-line using a silent installation

This installation method runs silently, so you do not see status messages during the actual installation.

About this task

Use the following steps to install the event synchronization using a silent installation from the command-line on your event server.

Procedure

1. Change to the `tec` directory on the IBM Tivoli Monitoring Tools DVD or DVD image.
2. Run the following command to generate the configuration file:

Windows

```
ESync3000Win32.exe -options-template filename
```

where *filename* is the name of the configuration file to create, for example, `es_silentinstall.conf`.

UNIX

```
ESync3000operating_system.bin -options-template filename
```

where *operating_system* is the operating system you are installing on (Aix, HP11, Linux, linux390, or Solaris). For example, run the following command on an AIX computer:

```
ESync3000Aix.bin -options-template filename
```

3. Edit the output file to specify the values described in Table 136 on page 747.

Note:

- a. Remove the number signs (###) from the beginning of any value that you want to specify.
- b. You must specify the following values:

```
### -P installLocation="value"
```

```
### -W configInfoPanel3.fileLocn="value"
```

and the following information for at least one monitoring server:

```
cmdSvrsPnlNotGuiMode.hostname1,  
cmdSvrsPnlNotGuiMode.userID1  
cmdSvrsPnlNotGuiMode.pswd1  
cmdSvrsPnlNotGuiMode.retypePswd1
```

If you do not specify any of the other values, the default values are used. If you specify values, ensure that the values you specify meets the minimum required values. Otherwise, the installation stops and an error is written to the log file.

Table 136. Netcool/OMNIBus event synchronization configuration values

Value	Description
installLocation	<p>The directory where you want to install the product. If the directory path contains spaces, enclose it in double quotation marks (" "). For example, to install the product to C:\Program Files\My Product, use</p> <pre>-P installLocation="C:\Program Files\My Product"</pre> <p>On Windows systems, the default installation directory is: C:\Program Files\IBM\SitForwarder</p> <p>On Linux/UNIX systems, the default installation directory is: /opt/IBM/SitForwarder</p>
configInfoPanel.filename	<p><i>For bidirectional architecture only:</i> The name of the file where IBM Tivoli Monitoring Situation Update Forwarder configuration information is stored. The default file name is situpdate.conf.</p>
configInfoPanel.pollingInt	<p><i>For bidirectional architecture only:</i> The polling interval, in seconds used by the IBM Tivoli Monitoring Situation Update Forwarder to determine if there are new event updates to send from Netcool/OMNIBus to the hub Tivoli Enterprise Monitoring Server. The minimum value is 1, while the default value is 3. If there are no situation events, the IBM Tivoli Monitoring Situation Update Forwarder rests for 3 seconds.</p>
configInfoPanel.crcByteCnt	<p><i>For bidirectional architecture only:</i> Number of bytes that the IBM Tivoli Monitoring Situation Update Forwarder uses when it saves the location of the last event it processes. This value must be an integer. The minimum (and default) value is 50.</p>
configInfoPanel.cmsSoapURL	<p><i>For bidirectional architecture only:</i> The URL for the SOAP Server configured on the computer where the hub monitoring server is running. The default value is cms/soap. Do not change this value. The actual URL for the SOAP Server is dynamically constructed from the incoming event.</p>
configInfoPanel.bufFlushRate	<p><i>For bidirectional architecture only:</i> The maximum number of event updates to be sent by the IBM Tivoli Monitoring Situation Update Forwarder to the hub monitoring server at one time. The minimum (and default) value is 10 events.</p>
configInfoPanel.logLevel	<p><i>For bidirectional architecture only:</i> The level of debug information for the IBM Tivoli Monitoring Situation Update Forwarder that is logged. You have the following choices:</p> <ul style="list-style-type: none"> • Low (default) • Medium • Verbose

Table 136. Netcool/OMNIBus event synchronization configuration values (continued)

Value	Description
configInfoPanel3.filesize	<i>For bidirectional architecture only:</i> The maximum permitted size, in bytes, for any one event cache file used by the IBM Tivoli Monitoring Situation Update Forwarder. The minimum (and default) value is 50000. Do not use commas when specifying this value (50,000 instead of 50000).
configInfoPanel3.fileNumber	<i>For bidirectional architecture only:</i> The maximum number of event caches files at any given time used by the IBM Tivoli Monitoring Situation Update Forwarder. The minimum value is 2, while the default value is 10. When this value is reached, the oldest file is deleted to make room for a new file.
configInfoPanel3.fileLocn	<i>For bidirectional architecture only:</i> The location where IBM Tivoli Monitoring Situation Update Forwarder event cache files are located. The default locations are as follows: <ul style="list-style-type: none"> • On Windows: C:\Program Files\IBM\SitForwarder\persistence. • On UNIX: /opt/IBM/SitForwarder/persistence
cmsSvrsPnlNotGuiMode.hostname# Note: The pound sign (#) stands for a number 1 - 10. For example, "hostname1".	<i>For bidirectional architecture only:</i> The host name of each hub monitoring server that sends events to the event server. For each hub monitoring server, specify one property with the monitoring server's fully qualified hostname and specify another property with the monitoring server's short hostname. Specify up to 10 monitoring server host names. If you have configured the Hot Standby feature, you must configure the host name information for both the primary and secondary hub monitoring servers.
cmsSvrsPnlNotGuiMode.userID#	<i>For bidirectional architecture only:</i> The user ID for the hub monitoring server, identified in host name#, to use to access the computer where the monitoring server is running.
cmsSvrsPnlNotGuiMode.pswd#	<i>For bidirectional architecture only:</i> The password for the user ID used to access the computer where the hub monitoring server is running.
cmsSvrsPnlNotGuiMode.retypePswd#	<i>For bidirectional architecture only:</i> The password confirmation for the user ID.

4. Save the file.

5. Run the following command:

Windows

```
ESync3000Win32.exe -options filename -silent
```

where *filename* is the name of your configuration file.

UNIX

```
ESync3000operating_system.bin -options filename -silent
```

where *operating_system* is the operating system you are installing on (Aix, HP11, Linux, linux390, Solaris). For example, on AIX, run the following command:

```
ESync3000Aix.bin -options filename -silent
```

Results

When installation is complete, the results are written to the `itm_tec_event_sync_install.log` file. On UNIX systems, this log file is always created in the `/tmp` directory. For Windows systems, this file is created in the directory defined by the `%TEMP%` environment variable. To determine where this directory is defined for the current command-line window, run the following command:

```
echo %TEMP%
```

If you specified the monitoring servers in the silent installation configuration file, you might consider deleting that file after installation, for security reasons. The passwords specified in the files are not encrypted.

If you want to define additional monitoring servers (in addition to the one required monitoring server), run the `sitconfuser.sh` command. Repeat this command for each monitoring server.

If you specified your monitoring servers after the installation, you must stop and restart the Situation Update Forwarder process manually.

Updating the Netcool/OMNIBus ObjectServer with IBM Tivoli Monitoring attributes, tables, and triggers

IBM Tivoli Monitoring provides updates to the Netcool/OMNIBus ObjectServer attributes, tables, and triggers. The type of updates required depend on whether you are using a single-tier architecture, a single-tier high-availability architecture, or a multitier architecture.

A single-tier architecture uses a single ObjectServer. A single-tier high-availability architecture uses a primary and backup ObjectServer with a bi-directional gateway between them. A multitier architecture uses multiple sets of ObjectServers. For more information on the last two architecture types, see “Netcool/OMNIBus Multitiered and High-availability Architecture” on page 728.

Table 137 indicates which ObjectServer updates are required for each architecture type.

Table 137. ObjectServer updates required for each architecture type

Task	single-tier architecture (single ObjectServer)	single-tier, high-availability architecture	Multitier architecture
“Updating gateways to map attributes” on page 750.		X Perform this task for the gateway between the primary and backup ObjectServers.	X Perform this task for each unidirectional and bidirectional gateway between the ObjectServers.
“Updating the bidirectional failover gateway to replicate tables” on page 751.		X Perform this task for the bidirectional gateway between the ObjectServers.	X Perform this task for each bidirectional gateway used between ObjectServers in the aggregation tier.
“Updating the OMNIBus database schema in the collection tier” on page 752.			X Perform this task on each ObjectServer in the collection tier.

Table 137. ObjectServer updates required for each architecture type (continued)

Task	single-tier architecture (single ObjectServer)	single-tier, high-availability architecture	Multitier architecture
“Updating the OMNIBus database schema on single-tier or aggregation tier ObjectServers” on page 753.	X	X Perform this task on each ObjectServer.	X Perform this task for each ObjectServer in the aggregation tier.
“Changing the default deduplication trigger” on page 755.	X	X Perform this task on each ObjectServer.	X Perform this task for each ObjectServer in the aggregation tier.
“Updating the OMNIBus database schema in the display tier” on page 758.			X Perform this task on each ObjectServer in the display tier.

Updating gateways to map attributes

If a Netcool/OMNIBus multitier architecture or single-tier high-availability architecture is used, then the gateways between ObjectServers need to be modified to add the relevant attributes for IBM Tivoli Monitoring integration with OMNIBus.

Before you begin

This task should be performed by a Netcool/OMNIBus administrator.

About this task

In a multitier architecture, update the map definition file for each of these gateways that exist in your environment:

- Collection tier to aggregation tier unidirectional gateways.
- Aggregation tier bidirectional failover gateways.
- Aggregation tier to display tier unidirectional gateways.

In a single-tier high-availability architecture, update the map definition file for the bidirectional failover gateway between the primary and backup ObjectServers.

Perform the following steps for each affected gateway.

Procedure

1. Copy the GATE_itm.map file from an ObjectServer where the IBM Tivoli Event Synchronization component was installed. The GATE_itm.map file can found be in the event_sync_install_dir/omnibus/multitier directory where event_sync_install_dir is the location where the IBM Tivoli Monitoring Event Synchronization component is installed.
2. Copy the mappings in the GATE_itm.map file into the gateway map definition file that also uses the .map extension. The following two types of mappings can be found in the GATE_itm.map file:
 - a. **Custom IBM Tivoli Monitoring attributes for the ObjectServer alerts.status table:** these attributes should be added to the STATUSMAP mapping entry in the gateway map definition file. Your gateway map file might contain this comment block to identify where to add these custom attributes:


```
#####
#
#      CUSTOM alerts.status FIELD MAPPINGS GO HERE
#
#####
```

- b. **Custom IBM Tivoli Monitoring table mappings:** Custom Tivoli Monitoring tables can be added to the ObjectServer database. The mapping definitions for custom tables only need to be copied to the gateway map definition file for bidirectional failover gateways:
 - in the aggregation tier of a multitier high-availability architecture, or
 - in a single-tier high-availability architecture

Copy the mapping definitions for the custom tables from the GATE_itm.map file to the end of the bidirectional failover gateway map definition file.

What to do next

You must restart the gateways for the changes to take effect. However, restart the gateways only after all other updates to the OMNIbus database schema are complete. See Table 132 on page 732 to determine when to restart the gateways.

Updating the bidirectional failover gateway to replicate tables

If a Netcool/OMNIbus multitier architecture or single-tier high-availability architecture is used, then the bidirectional gateways between ObjectServers must be modified to replicate IBM Tivoli Monitoring custom tables in the ObjectServer database.

Before you begin

This task should be performed by a Netcool/OMNIbus administrator.

About this task

In a multitier high-availability architecture, update the bidirectional failover gateways used between the ObjectServers in the aggregation tier. In a single-tier high-availability architecture, update the bidirectional failover gateway between the primary and backup ObjectServers.

Perform the following steps for each affected bidirectional failover gateway.

Procedure

1. Copy the GATE_itm.tblrep.def file from an ObjectServer where the IBM Tivoli Event Synchronization component was installed. You can find the GATE_itm.tblrep.def file in the event_sync_install_dir/omnibus/multitier directory where event_sync_install_dir is the location where the IBM Tivoli Monitoring Event Synchronization component is installed.
2. Copy the table replication definitions in the GATE_itm.tblrep.def file to the bottom of the gateway table replication definition file that also uses the .tblrep.def file extension. Use the following criteria to determine which table definitions to copy:
 - a. Always copy the replication definition for the alerts.itm_loopback_events table.
 - b. Copy the replication definition for the alerts.itm_heartbeat_events table if monitoring agents will be sending events directly to Netcool/OMNIbus and the agents will be configured to send heartbeat events. If you plan to use agent heartbeat settings, you must also run the itm_heartbeat.sql file when updating the OMNIbus database schema in a single-tier architecture or in the aggregation tier of a multitier architecture.
 - c. Copy the replication definition for the alerts.itm_cleared_event_cache table if you are planning to use bidirectional event synchronization between the monitoring servers and ObjectServer. If you

copy this definition, you must also execute the `itm_event_cache.sql` file when updating the OMNIBus database schema in a single-tier architecture or in the aggregation tier of a multitier architecture.

What to do next

You must restart the gateways for the changes to take effect. However, the restart should occur after all other updates to the OMNIBus database schema are complete. See Table 132 on page 732 to determine when to restart the gateways.

Updating the OMNIBus database schema in the collection tier

In a Netcool/OMNIBus multitier architecture, add IBM Tivoli Monitoring triggers to each ObjectServer in the collection tier.

Before you begin

This task should be performed by a Netcool/OMNIBus administrator.

About this task

Perform the following steps for each Netcool/OMNIBus Object Server defined in the collection tier.

Procedure

1. Copy the `collection_itm.sql` file from an ObjectServer in the aggregation tier to each ObjectServer in the collection tier. The `collection_itm.sql` file can be found in the `event_sync_install_dir/omnibus/multitier` directory where `event_sync_install_dir` is the location where the IBM Tivoli Monitoring Event Synchronization component is installed on an ObjectServer in the aggregation tier.
2. Update the Object Server database with the following command, which pipes the SQL command set into the SQL command-line tool and performs the updates to the ObjectServer database:

- **Windows**

```
type path_to_file\collection_itm.sql | %OMNIHOME%\..\bin\redist\isql
-U username
-P password
-S server_name
```

where:

\$OMNIHOME

Is the system-defined variable defining the installation location of OMNIBus.

username

Is the OMNIBus Object Server user name.

password

Is the OMNIBus Object Server password.

server_name

Is the OMNIBus Object Server name defined for process control.

path_to_file

Is the fully qualified path to the specified SQL file.

- **UNIX**

```
$OMNIHOME/bin/nco_sql -user username
-p password password
-s server server_name
< path_to_file/collection_itm.sql
```

where:

\$OMNIHOME

Is the system-defined variable defining the installation location of OMNIBus

username

Is the OMNIBus ObjectServer user name

password

Is the OMNIBus ObjectServer password

server_name

Is the OMNIBus ObjectServer name defined for process control.

path_to_file

Is the fully qualified path to the specified SQL file.

Updating the OMNIBus database schema on single-tier or aggregation tier ObjectServers

Update the Netcool/OMNIBus ObjectServer with the attributes and triggers that allow the ObjectServer to correctly process IBM Tivoli Monitoring events and forward events to the IBM Tivoli Monitoring Situation Update Forwarder when using bidirectional communication.

To verify if specific customization settings are needed, see “Customizing how the IBM Tivoli Monitoring OMNIBus triggers handle event status updates from the monitoring servers” on page 775.

Perform the steps in this section for each ObjectServer in a single-tier architecture or for each ObjectServer in the aggregation tier of a multitier architecture.

The tasks in this section should be completed by the Netcool/OMNIBus Administrator.

If you are setting up a unidirectional architecture and have installed the IBM Tivoli Monitoring Event Synchronization Component on a system other than your Netcool/OMNIBus ObjectServer, you must manually copy the SQL files to your ObjectServer. The files are located in the `event_sync_install_dir/omnibus` directory. Also, for a unidirectional architecture, you must modify the contents of the `itm_proc.sql` file, changing every instance of the word *REJECT* to the word *ACCEPT* and ensure that the path specified for the `itmfile` variable exists on the system with Netcool/OMNIBus ObjectServer. You only need to load the `itm_proc.sql` and `itm_db_update.sql` files. If you later decide to convert your environment to a bidirectional architecture, you must also load the `itm_sync.sql` file.

If you are setting up a bidirectional architecture, you might want to customize the `itm_proc.sql` file before loading it. For more information, see “Customizing how the IBM Tivoli Monitoring OMNIBus triggers handle event status updates from the monitoring servers” on page 775. Also review the information at the end of the topic about ensuring that the `eventcmd` procedure is correctly configured to send events back to the hub monitoring server.

Note: If you are integrating IBM Tivoli Monitoring and Tivoli Business Service Manager, you must also complete the following steps:

1. If you are using Tivoli Business Service Manager V4.2.1, add the Tivoli Business Service Manager schema updates before you add the Tivoli Monitoring schema updates. If you add the Tivoli Monitoring schema updates before the Tivoli Business Service Manager schema updates, rerun the procedure in this section to add the IBM Tivoli Monitoring schema updates again to ensure that the latest definitions are used.
2. After updating the OMNIBus schema with the Tivoli Monitoring updates, run the Tivoli Business Service Manager discover schema utility (`rad_discover_schema`). See the Tivoli Business Service Manager Information Center for detailed instructions on using this utility: <http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/index.jsp?topic=/com.ibm.tivoli.itbsm.doc/welcome.htm>.

3. After running the discover schema utility, remember to restart the Tivoli Business Service Manager Dataserver. Failure to do so can cause connection problems.

If the OMNIBus ObjectServer is running on a UNIX system as a non-root user and the event synchronization component is installed and run as either root or another user, verify that the user under which the ObjectServer is running has write permission to the `event_sync_install_dir/log` directory before updating the OMNIBus database schema. If the user does not have write permission to this directory there will be error messages related to `itmfile` when the `itm_proc.sql` file is loaded into Netcool/OMNIBus and ITM-OMNIBus event synchronization will not work.

The command to configure OMNIBus pipes the SQL command set into the SQL command-line tool and performs the updates to the ObjectServer.

1. Update the Object Server database with the following commands:

- **Windows**

```
type path_to_file\itm_proc.sql | %OMNIHOME%\..\bin\redist\isql -U username
-P password
-S server_name
```

```
type path_to_file\itm_db_update.sql | %OMNIHOME%\..\bin\redist\isql -U username
-P password
-S server_name
```

Only run this command if you are using the bidirectional architecture:

```
type path_to_file\itm_sync.sql | %OMNIHOME%\..\bin\redist\isql -U username
-P password
-S server_name
```

Only run this command if you are using the bidirectional architecture:

```
type path_to_file\itm_event_cache.sql | %OMNIHOME%\..\bin\redist\isql -U username
-P password
-S server_name
```

where:

\$OMNIHOME

Is the system-defined variable defining the installation location of OMNIBus.

username

Is the OMNIBus Object Server user name.

password

Is the OMNIBus Object Server password.

server_name

Is the OMNIBus Object Server name defined for process control.

path_to_file

Is the fully qualified path to the specified SQL file.

- **UNIX**

```
$OMNIHOME/bin/nco_sql -user username
-p password password
-s server server_name
< path_to_file/itm_proc.sql
```

```
$OMNIHOME/bin/nco_sql -user username
-p password password
-s server server_name
< path_to_file/itm_db_update.sql
```

Only run this command if you are using the bidirectional architecture:

```
$OMNIHOME/bin/nco_sql -user username  
-password password  
-server server_name  
< path_to_file/itm_sync.sql
```

Only run this command if you are using the bidirectional architecture:

```
$OMNIHOME/bin/nco_sql -user username  
-password password  
-server server_name  
< path_to_file/itm_event_cache.sql
```

where:

\$OMNIHOME

Is the system-defined variable defining the installation location of OMNIBus.

username

Is the OMNIBus ObjectServer user name.

password

Is the OMNIBus ObjectServer password.

server_name

Is the OMNIBus ObjectServer name defined for process control.

path_to_file

Is the fully qualified path to the specified SQL file. The SQL files are located in the event_sync_install_dir/omnibus directory where event_sync_install_dir is the location where the IBM Tivoli Monitoring Event Synchronization component is installed.

Note:

1. "Object exists" and "Attempt to insert duplicate row" errors occur if the scripts were previously run. These errors are harmless.
2. If you are running the SQL database schema for the first time, the following harmless error messages are displayed:
 - ERROR=Object not found on line 4 of statement "-- A workspace table for the ITM event clear automation...", at or near "itm_event_clear".
 - ERROR=Object not found on line 1 of statement "delete from alerts.itm_problem_events;...", at or near "itm_problem_events".
 - ERROR=Object not found on line 1 of statement "drop table alerts.itm_problem_events;...", at or near "itm_problem_events".

Changing the default deduplication trigger

The deduplication triggers provided by the IBM Tivoli Monitoring event synchronization component and the standard deduplication triggers provided with Netcool/OMNIBus are in conflict with each other. The standard deduplication triggers modify the Severity and Summary attributes that results in those attributes having incorrect values for IBM Tivoli Monitoring events for both unidirectional and bidirectional event synchronization.

The standard deduplication triggers must be modified by the Netcool/OMNIBus administrator to ignore events from IBM Tivoli Monitoring. The deduplication trigger to modify depends on the Netcool/OMNIBus architecture type:

single-tier architecture:

- In a single-tier architecture without high availability, modify the deduplication trigger.
- In a single-tier architecture with a bidirectional failover gateway for high availability, modify the agg_deduplication trigger in each ObjectServer.

Multitier architecture:

- In a multitier architecture without high availability, modify the deduplication trigger on the Objectserver in the aggregation tier.
- In a multitier architecture with a bidirectional failover gateway for high availability, modify the agg_deduplication trigger on each ObjectServer in the aggregation tier.

You can modify the appropriate trigger from either the Netcool/OMNIBus Administrator nco_config interface or by using the nco_sql CLI on the Netcool/OMNIBus ObjectServer.

Changing the default deduplication trigger using Netcool/OMNIBus Administrator:

This section outlines the steps for changing the default deduplication trigger using Netcool/OMNIBus Administrator.

Procedure

1. Run the nco_config utility to start Netcool/OMNIBus Administrator.
2. Connect to the Netcool/OMNIBus ObjectServer where the change will be made.
3. Select **AutomationTriggers** from the menu.
4. Go to the editor for the trigger to be modified. See the preceding paragraphs to determine if you should modify the deduplication or agg_deduplication trigger.
5. On the **When** tab, enter: **new.Type != 20** and **new.Type != 21**.
6. Save the trigger and exit from nco_config interface.

Using the nco_sql command to modify the deduplication trigger in environments without a bidirectional failover gateway:

This section shows you how to use the nco_sql command to modify the deduplication trigger in environments without a bidirectional failover gateway.

Procedure

1. Open the following file:
 - **Windows** %OMNIHOME%\etc\automation.sql.
 - **Linux** **UNIX** \$OMNIHOME/etc/automation.sql.
2. Locate the command that creates the deduplication trigger. See the following example of the deduplication trigger command:

Note: Your deduplication trigger might be different than what is shown in the following example depending on the Netcool/OMNIBus release that you are using or if you have customized the trigger.

```
create or replace trigger deduplication
group default_triggers
priority 1
comment 'Deduplication processing for ALERTS.STATUS'
before reinsert on alerts.status
for each row
begin
  set old.Tally = old.Tally + 1;
  set old.LastOccurrence = new.LastOccurrence;
  set old.StateChange = getdate();
  set old.InternalLast = getdate();
  set old.Summary = new.Summary;
  set old.AlertKey = new.AlertKey;
  if (( old.Severity = 0) and (new.Severity > 0))
  then
```



```

    set old.Severity = new.Severity;
  end if;
end;
go

```

- Copy this command to a temporary file. In this example, the file name is /tmp/dedup.sql.
- Edit /tmp/dedup.sql and add the line that is highlighted in bold in the following example command so that the trigger ignores IBM Tivoli Monitoring events:

```

create or replace trigger deduplication
group default_triggers
priority 1
comment 'Deduplication processing for ALERTS.STATUS'
before reinsert on alerts.status
for each row
when (new.Type != 20) and (new.Type != 21)
begin
  set old.Tally = old.Tally + 1;
  set old.LastOccurrence = new.LastOccurrence;
  set old.StateChange = getdate();
  set old.InternalLast = getdate();
  set old.Summary = new.Summary;
  set old.AlertKey = new.AlertKey;
  if (( old.Severity = 0) and (new.Severity > 0))
  then
    set old.Severity = new.Severity;
  end if;
end;
go

```

- Save the file, and run the following command to replace the deduplication trigger:
 - Windows** %OMNIHOME%\..\bin\redist\isql -U username -P password -S server_name < C:\tmp\dedup.sql.
 - Linux** **UNIX** \$OMNIHOME/bin/nco_sql -user username -password password -server server_name < /tmp/dedup.sql.

where:

OMNIHOME

Is the system-defined variable that defines the installation location of Netcool/OMNIBus ObjectServer.

username

Is the Netcool/OMNIBus ObjectServer user name.

password

Is the Netcool/OMNIBus ObjectServer password.

server_name

Is the Netcool/OMNIBus ObjectServer name defined for process control.

Using nco_sql to modify the agg_deduplication trigger in environments with a bidirectional failover gateway:

This section shows you how to use the nco_sql command to modify the deduplication trigger in environments with a bidirectional failover gateway.

Procedure

- Open the following file:
 - Windows** %OMNIHOME%\extensions\multitier\objectserver\aggregation.sql.
 - Linux** **UNIX** \$OMNIHOME/extensions/multitier/objectserver/aggregation.sql.
- Copy the agg_deduplication trigger logic to a temporary file. In this example, the temporary file name is /tmp/agg_dedup.sql.

3. Edit /tmp/agg_dedup.sql and add the line highlighted in bold in the following example so that the trigger ignores IBM Tivoli Monitoring events.

Note: Only the beginning portion of the agg_deduplication trigger is shown in the example that follows. The remaining logic is not shown because it does not need any modifications. Your agg_deduplication trigger might be different than what is shown in the following example depending on the Netcool/OMNIBus release that you are using or if you have customized the trigger.

```
CREATE OR REPLACE TRIGGER agg_deduplication
GROUP default_triggers
PRIORITY 2
COMMENT 'Replacement reinsert trigger (alerts.status) for multi-ObjectServer environments.'
BEFORE REINSERT ON alerts.status
FOR EACH ROW
when (new.Type != 20) and (new.Type != 21)
declare

    now utc;
begin
```

4. Save the file, and run the following command to replace the agg_deduplication trigger:

- **Windows** %OMNIHOME%\..\bin\redist\isql -U username -P password -S server_name < C:\tmp\agg_dedup.sql.
- **Linux** **UNIX** \$OMNIHOME/bin/nco_sql -user username -password password -server server_name < /tmp/agg_dedup.sql.

where:

OMNIHOME

Is the system-defined variable that defines the installation location of Netcool/OMNIBus ObjectServer.

username

Is the Netcool/OMNIBus ObjectServer user name.

password

Is the Netcool/OMNIBus ObjectServer password.

server_name

Is the Netcool/OMNIBus ObjectServer name defined for process control.

Updating the OMNIBus database schema in the display tier

In a Netcool/OMNIBus multitier architecture, add IBM Tivoli Monitoring attributes to the ObjectServer alerts.status table in the display tier.

Before you begin

This task should be performed by a Netcool/OMNIBus administrator.

About this task

For each Netcool/OMNIBus Object Server defined in the display tier, complete the following steps:

Procedure

1. Copy the display_itm.sql file from an ObjectServer in the aggregation tier to each ObjectServer in the display tier.

The display_itm.sql file can be found in the event_sync_install_dir/omnibus/multitier directory where event_sync_install_dir is the location where the IBM Tivoli Monitoring Event Synchronization component is installed on an ObjectServer in the aggregation tier.

2. Update the Object Server database with the following command, which pipes the SQL command set into the SQL command-line tool and performs the updates to the ObjectServer database:

- Windows

```
type path_to_file\display_itm.sql | %OMNIHOME%\..\bin\redist\isql
-U username
-P password
-S server_name
```

where:

\$OMNIHOME

Is the system-defined variable defining the installation location of OMNIbus.

username

Is the OMNIbus Object Server user name.

password

Is the OMNIbus Object Server password.

server_name

Is the OMNIbus Object Server name defined for process control.

path_to_file

Is the fully qualified path to the specified SQL file.

- UNIX

```
$OMNIHOME/bin/nco_sql -user username
-p password password
-s server server_name
< path_to_file/collection_itm.sql
```

where:

\$OMNIHOME

Is the system-defined variable defining the installation location of OMNIbus

username

Is the OMNIbus ObjectServer user name

password

Is the OMNIbus ObjectServer password

server_name

Is the OMNIbus ObjectServer name defined for process control.

path_to_file

Is the fully qualified path to the specified SQL file.

Starting and stopping the IBM Tivoli Monitoring Situation Update Forwarder

If you are using bidirectional architecture to forward event status updates from Netcool/OMNIbus back to the hub monitoring server, you must start the IBM Tivoli Monitoring Situation Update Forwarder.

The tasks in this section should be completed by the Netcool/OMNIbus Administrator.

This process is started automatically when the Netcool/OMNIbus ObjectServer starts. To start the process manually, change to the `event_sync_install_dir/bin` directory where `event_sync_install_dir` is the directory where the IBM Tivoli Monitoring Event Synchronization Component is installed. Then run the following command:

- Windows

```
startSUF.cmd
```

UNIX

```
startSUF.sh
```

To stop the process, run the following command:

Windows

```
stopSUF.cmd
```

UNIX

```
stopSUF.sh
```

On Windows, you can also start and stop the IBM Tivoli Monitoring Situation Update Forwarder service to start or stop the forwarding of event updates. You can start and stop this service either from the Windows Service Manager utility or with the following commands:

```
net start situpdate
net stop situpdate
```

Configuring the Netcool/OMNIBus EIF probe

Use this step to configure the Netcool/OMNIBus EIF Probe with the rules for mapping situation EIF events to OMNIBus events.

The default mapping of IBM Tivoli Monitoring EIF attribute slots to OMNIBus attributes is shown in “Default mapping of situation events to OMNIBus events” on page 787. To configure the mapping, you must uncomment the include statement for the `itm_event.rules` in the probe's `tivoli_eif.rules` file and copy the IBM Tivoli Monitoring rules files to the probe's operating system directory.

Three Tivoli Monitoring rules files are available:

- The `itm_event.rules` file is included by the `tivoli_eif.rules` probe rules file. This is located in the omnibus directory under the event synchronization component's installation directory.
- The `tbsm_eif_event.rules` file is included by the `itm_event.rules` file and sets attributes specific to Tivoli Business Service Manager integration. This is located in the omnibus/tbsm directory under the event synchronization component's installation directory.
- The `itm_custom_overrides.rules` is included by the `itm_event.rules` file and can be used to customize the IBM Tivoli Monitoring rules behavior. See “Customizing the rules file” on page 762 for more details. This is located in the omnibus directory under the event synchronization component's installation directory.

If IBM Tivoli Monitoring is sending virtualization events or predictive analytics events to Netcool/OMNIBus, there are rules and triggers available with Netcool/OMNIBus ObjectServer for additional event enrichment and correlation handling. These rules also set some of the OMNIBus attributes (including Node and Summary) differently from the values that are normally used for monitoring events. See the Netcool/OMNIBus Information Center for details on extending the functionality of Netcool/OMNIBus for virtualization and predictive events. The virtualization rules and predictive rules files have include statements in the `tivoli_eif.rules` file. Verify that you are using a version of the EIF Probe's `tivoli_eif.rules` file that includes the virtualization rules file (`tivoli_eif_virtualization_pt2.rules`) and predictive rules file (`predictive_event.rules`) after `itm_event.rules`. If either of these rules files are included before `itm_event.rules` then move the include statements. You should also verify that you are using Netcool/OMNIBus Version 7.3.0 FP 6 or later, or Netcool/OMNIBus Version 7.3.1 FP2 or later, as those fix packs have the required updates to the rules files for integration with IBM Tivoli Monitoring.

Some monitoring agents might also provide their own probe rules include files that should also be included after `itm_event.rules` in the probe's `tivoli_eif.rules` file. See the documentation for the specific agents you are using in your environment to determine if the agent has its own rules file.

Note: You must recycle the probe after you update any of the rules files.

If you are using a Netcool/OMNIbus multitier architecture, you should also ensure the EIF Probe is configured to connect to an ObjectServer in the collection tier.

Updating the rules files of the EIF probe

Perform the steps in this section to update the rules files of the EIF probe.

Before you begin

The tasks in this section should be completed by the Netcool/OMNIbus Administrator.

Procedure

1. Uncomment the include statement for `itm_event.rules` in the `tivoli_eif.rules` file in the probe's operating system directory:

`%OMNIHOME%\probes\arch` (Windows)

or

`$OMNIHOME/probes/arch` (UNIX)

where:

OMNIHOME

Is a system-defined variable defining the installation location of OMNIbus.

arch

Represents the operating system directory on which the probe is installed; for example, `solaris2` when running on a Solaris system, and `win32` for a Windows system.

Note: If the probe's version of the `tivoli_eif.rules` does not contain a statement to include the `itm_event.rules` file, you must upgrade to version 10 or later of the EIF Probe.

2. Copy the `itm_event.rules` file in the omnibus directory where the event synchronization component is installed, to this directory on the machine where the EIF probe is installed:

`%OMNIHOME%\probes\arch` (Windows)

or

`$OMNIHOME/probes/arch` (UNIX)

where:

OMNIHOME

Is a system-defined variable defining the installation location of OMNIbus.

arch

Represents the operating system directory on which the probe is installed; for example, `solaris2` when running on a Solaris system, and `win32` for a Windows system.

3. If you are integrating Tivoli Business Service Manager as well, uncomment the line in `itm_event.rules` with the include statement for `tbsm_eif_event.rules`, and copy the `tbsm_eif_events.rules` from the omnibus/tbsm directory to this directory on the machine where the EIF probe is installed:

`%OMNIHOME%\probes\arch` (Windows)

or

`$OMNIHOME/probes/arch` (UNIX)

where:

OMNIHOME

Is a system-defined variable defining the installation location of OMNIBus.

arch

Represents the operating system directory on which the probe is installed; for example, solaris2 when running on a Solaris system, and win32 for a Windows system.

4. If rules customizations are done in `itm_custom_override.rules`, uncomment the line in `itm_event.rules` with the include statement for `itm_custom_override.rules`. Copy the `itm_custom_override.rules` file with the customizations to this directory on the machine where the EIF probe is installed:

`%OMNIHOME%\probes\arch` (Windows)

or

`$OMNIHOME/probes/arch` (UNIX)

where:

OMNIHOME

Is a system-defined variable defining the installation location of OMNIBus.

arch

Represents the operating system directory on which the probe is installed; for example, solaris2 when running on a Solaris system, and win32 for a Windows system.

5. Stop the EIF probe.
 - **Windows** In the Control Panel, open Administrative Tools, then Services. In the list of services, double-click the EIF probe and then click Stop.

- **UNIX** Issue the following command:

`$OMNIHOME/bin/nco_pa_stop -process probe_name`

where:

\$OMNIHOME

Is the system-defined variable defining the installation location of OMNIBus.

probe_name

Is the OMNIBus EIF probe name defined for Process Control.

6. Restart the OMNIBus probe.
 - **Windows** In the list of services, double-click OMNIBus EIF Probe and then click **Start**.
 - **UNIX** Issue the following command:
`$OMNIHOME/bin/nco_pa_start -process probe_name`

Customizing the rules file

To customize the IBM Tivoli Monitoring probe rules, add your customizations to the `itm_custom_override.rules` file.

The `itm_custom_override.rules` file can be included by `itm_event.rules` by uncommenting the line with the include statement for it. You should not modify any of the OMNIBus attributes that are used by the IBM Tivoli Monitoring triggers to perform event synchronization. See “Default mapping of situation events to OMNIBus events” on page 787 for the list of attributes you should not change.

If you want to map agent specific slots to Netcool/OMNIBus attributes, see “Generic mapping for agent-specific slots” on page 792 for information on slot naming and other considerations.

The `itm_custom_overrides.rules` file is copied into the following directory:

`%OMNIHOME%\probes\arch` (Windows)

or

`$OMNIHOME/probes/arch` (UNIX)

where:

OMNIHOME

Is a system-defined variable defining the installation location of OMNIBus.

arch

Represents the operating system directory on which the probe is installed; for example, `solaris2` when running on a Solaris system, and `win32` for a Windows system.

After updating the rules file, restart the probe using the instructions in “Configuring the Netcool/OMNIBus EIF probe” on page 760.

Configuring the OMNIBus server for program execution from scripts

To run the IBM Tivoli Monitoring Situation Update Forwarder program from SQL automation scripts for sending synchronization events to the hub monitoring server, the Netcool/OMNIBus ObjectServer must be running under process control. You must also set the **PA.Username** and **PA.Password** properties in `OMNIHOME/etc/NCOMS.props` file, where `OMNIHOME` is the system-defined variable defining the installation location of Netcool/OMNIBus.

Before you begin

The tasks in this section should be completed by the Netcool/OMNIBus Administrator if you are using bidirectional event synchronization. If you are using a Netcool/OMNIBus multitier architecture, these tasks should only be performed on the ObjectServers in the aggregation tier and are not necessary for the ObjectServers in the collection and display tiers.

About this task

For Linux and UNIX: The **PA.Username** property must be set to the user name for connecting to the process control agent. The **PA.Password** property must be set to the password of the user connecting to the process control agent. By default, process control grants access to the members of `ncoadmin` group so configure **PA.Username** to specify a user in this group.

For Windows: The **PA.Username** property must be set to a Windows account name, and the **PA.Password** property must be set to the password for that account.

See OMNIBus documentation for more information about configuring OMNIBus server under process control and for information about the `nco_pa_crypt` utility that encrypts the **PA.Password** property value.

After you change the **PA.Username** and **PA.Password** properties in the `OMNIHOME/etc/NCOMS.props` file, complete the following procedure to restart the OMNIBus ObjectServer:

Procedure

1. Stop the OMNIBus server:

- **Windows** In the Control Panel, open Administrative Tools, then Services. In the list of services, double-click OMNIBus server, then click Stop.
- **UNIX** Issue the following command from command-line
`$OMNIHOME/bin/nco_pa_stop -process server_name`

where:

\$OMNIHOME

Is the system-defined variable defining the installation location of OMNIBus.

server_name

Is the OMNIBus ObjectServer name defined for process control.

2. Restart the OMNIBus server.

- **Windows** In the list of services, double-click OMNIBus server, then click Start.
- **UNIX** Issue the following command from command-line:

```
$OMNIHOME/bin/nco_pa_start -process server_name
```

Configuring the hub monitoring server to forward events

Before the hub monitoring server forwards any situation events to Netcool/OMNIBus, you must enable event forwarding and define a default event destination.

The tasks in this section should be completed by the IBM Tivoli Monitoring Administrator.

Take the following steps to enable event forwarding for each hub monitoring server from which you want to forward events to an OMNIBus ObjectServer. If the Hot Standby feature is configured, you should perform these steps for both the primary and secondary hub monitoring servers:

Windows **monitoring servers only**, perform the following steps:

1. Open Manage Tivoli Enterprise Monitoring Services.
2. Right-click the monitoring server and click **Reconfigure**.
3. On the configuration options window, select **Tivoli Event Integration Facility**.
4. Click **OK** and **OK**.
5. Complete the following fields on the Event Server: Location and Port Number window and click **OK**:

Server or EIF Probe Location

Type the host name or IP address for the computer where the EIF probe is installed.

Port Number

Type the port number on which the probe is listening. Check with the Netcool/OMNIBus administrator that installed the EIF probe to determine what port number the EIF probe is using, or see the *IBM Tivoli Monitoring Troubleshooting Guide* for instructions on how to determine the EIF probe port number. If the default probe port number was used, the number depends on the version of the probe and how the probe was installed. If you are using Netcool/OMNIBus Probe for Tivoli EIF Version 8 or later, the default port number is 9998. If Tivoli Business Service Manager V4.2.1 was used to install the probe, the default probe port number is 5530.

Linux **UNIX** **monitoring servers:** You configured the EIF probe and port information for the Linux/UNIX monitoring server during installation, if you installed the monitoring server using the configuration instructions in this installation guide. However, if you did not configure this information, see “Configuring the hub monitoring server” on page 202 for the procedure.

You can use the `tacmd createEventDest` command to configure additional event destinations if the hub monitoring server should forward events to more than one Netcool/OMNIBus Probe for EIF. For more details on this command, see the *IBM Tivoli Monitoring Command Reference*. For details on how to specify which event destinations to use for each situation, see “Specifying which situations forward events to Netcool/OMNIBus” on page 780.

Controlling event forwarding rate

The `om_tec.config` file controls the forwarding of events from IBM Tivoli Monitoring to IBM Tivoli Netcool/OMNIBus.

This file contains the following parameter:

BufferFlushRate=events_per_minute

Specifies the number of events that are sent per minute when the adapter has reestablished its connection to the Netcool/OMNIBus Probe for Tivoli EIF. Once the adapter has recovered the lost connection, if there are events in the buffer, the events are sent at this rate per minute. The default value is 0, meaning all events are sent in one burst.

If your environment can have many open situation events, you might want to adjust this parameter to control the rate at which events are sent to the probe. To edit this file and change this parameter:

- **Windows**
 1. Open Manage Tivoli Enterprise Monitoring Services.
 2. Right-click Tivoli Enterprise Monitoring Server, and click Advanced → Edit EIF Configuration.
- **Linux** **UNIX** Edit file *install_dir/tables/hostname/TECLIB/om_tec.config*, where *install_dir* is your Tivoli Monitoring installation directory and *hostname* is the name of the host running this monitoring server.

Restart the hub monitoring server after completing these steps:

- **Windows**
 1. Start the Manage Tivoli Enterprise Monitoring Services utility (**Start → (All) Programs → IBM Tivoli Monitoring → Manage Tivoli Monitoring Services**).
 2. Right-click the hub monitoring server and select **Recycle** from the popup menu.
- **Linux** **UNIX** Enter the following commands:

```
$ITM_INSTALL_PATH/bin/itmcmd stop tems_server
$ITM_INSTALL_PATH/bin/itmcmd start tems_server
```

where:

\$ITM_INSTALL_PATH

Is the system variable defining the installation location of the monitoring server.

tems_server

Is the host name of the computer where the monitoring server is installed (the host where you are executing this command).

Verifying installation and configuration

After following the steps in the previous section, your Netcool/OMNIBus Probe for Tivoli EIF and ObjectServer should be receiving IBM Tivoli Monitoring situation events. Before additional customization, you should verify that the event synchronization component is installed and event integration is configured successfully.

The verification requires coordination between the Netcool/OMNIBus and IBM Tivoli Monitoring administrators.

To verify unidirectional event flow:

Complete the steps in this section to verify unidirectional event flow.

Procedure

1. Start the Tivoli Enterprise Portal.
2. Ensure you have a situation that has been configured to forward EIF events, and that the situation condition can be easily created.
3. After you have created a situation condition, use the Tivoli Enterprise Portal to verify that one or more situation events are open. The Tivoli Enterprise Portal Physical Navigator displays an event severity indicator at the top level of the Navigator and on all affected nested nodes in the view. As a result, events generated by a situation are always visible. You can also see the complete list of open situation events in the Situation Event Console workspace.

4. Start the Netcool/OMNIBus Event List UI. For example, on the OMNIBus server machine, run the following command to start the Netcool/OMNIBus Native Event List: `$OMNIHOME/bin/nco_event` where `$OMNIHOME` is the directory location of the Netcool/OMNIBus installation.
5. Refresh the event view in the Netcool/OMNIBus Event List UI and check for situation events from IBM Tivoli Monitoring. You should see the same set of events that are in the Tivoli Enterprise Portal Situation Event Console.

Results

If you see events from IBM Tivoli Monitoring in the OMNIBus Event List UI, you have successfully set up unidirectional event flow in your environment.

If you do not see events from IBM Tivoli Monitoring in the OMNIBus Event List UI, see the *IBM Tivoli Monitoring Troubleshooting Guide* for information about how to resolve any issues.

To verify bidirectional event flow:

Complete the steps in this section to verify bidirectional event flow.

Procedure

1. Follow the steps in “To verify unidirectional event flow:” on page 765 to verify unidirectional event flow in your environment.
2. In the event view in the Netcool/OMNIBus Event List UI, perform an action (acknowledge, deacknowledge, delete, or clear) on one of the forwarded situation events from IBM Tivoli Monitoring.
3. In the Tivoli Enterprise Portal, verify that the state of the situation event changes to the expected state. See Table 125 on page 700 for details on the expected behavior of events.

Results

If you see the status of the corresponding situation change within the Tivoli Enterprise Portal, you have successfully set up bidirectional event flows in your environment.

If you do not see the status of the corresponding situation event change within the Tivoli Enterprise Portal, see the *IBM Tivoli Monitoring Troubleshooting Guide* for information about resolving issues.

Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIBus integration

This section contains the tasks you must perform if you are upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIBus integration.

If you have already installed and configured event integration between IBM Tivoli Monitoring and Netcool/OMNIBus and you are upgrading IBM Tivoli Monitoring, perform the tasks listed in Table 138 on page 767 to upgrade each of the environments that will be receiving events from the latest version of IBM Tivoli Monitoring.

Table 138. Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIbus integration.

	Task	Architecture Type	Administrator
1.	<p>Upgrade Netcool/OMNIbus ObjectServer, if necessary, to a release or fix pack version required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIbus Information Center for instructions on performing the upgrade: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIbus” on page 91.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
2.	<p>Back up the rules files in the Netcool/OMNIbus Probe for Tivoli EIF's operating system directory and then upgrade the Netcool/OMNIbus Probe for Tivoli EIF, if necessary, to the release required by IBM Tivoli Monitoring.</p> <p>See the Netcool/OMNIbus Information Center for instructions on performing the upgrade: http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm. See also “Required software for event integration with Netcool/OMNIbus” on page 91.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
3.	“Upgrading the IBM Tivoli Monitoring Event Synchronization Component” on page 768.	Unidirectional and bidirectional	Netcool/OMNIbus or IBM Tivoli Monitoring
4.	“Upgrading existing IBM Tivoli Monitoring events to use new OMNIbus attributes” on page 769.	Unidirectional and bidirectional	Netcool/OMNIbus

Table 138. Upgrading from a previous installation of IBM Tivoli Monitoring and Netcool/OMNIbus integration. (continued)

	Task	Architecture Type	Administrator
5.	<p>“Updating the Netcool/OMNIbus ObjectServer with IBM Tivoli Monitoring attributes, tables, and triggers” on page 749.</p> <p>Note: If you have made any customizations to the IBM Tivoli Monitoring triggers in the Netcool/OMNIbus ObjectServer, you should make the same updates to the new version of the triggers and procedures. You can edit the .sql files provided with the IBM Tivoli Monitoring Event Synchronization before loading them into Netcool/OMNIbus ObjectServer. For a list of customizations that might have been made to the existing triggers, see “Customizing how the IBM Tivoli Monitoring OMNIbus triggers handle event status updates from the monitoring servers” on page 775.</p>	Unidirectional and bidirectional	Netcool/OMNIbus
6.	Ensure that you have a default deduplication trigger. For more information, see “Replacing the default deduplication trigger” on page 770.	Unidirectional and bidirectional	Netcool/OMNIbus
7.	If you did not replace the default deduplication trigger in the previous task, ensure that you have changed the existing default duplication trigger to ignore IBM Tivoli Monitoring events. See “Changing the default deduplication trigger” on page 755 for details.	Unidirectional and bidirectional	Netcool/OMNIbus
8.	“Updating the EIF Probe Rules” on page 771.	Unidirectional and bidirectional	Netcool/OMNIbus

Upgrading the IBM Tivoli Monitoring Event Synchronization Component

When you upgrade the IBM Tivoli Monitoring event synchronization component for IBM Tivoli Netcool/OMNIbus you install the newest version of the Situation Update Forwarder, and replace any scripts that have been updated since the base version.

The latest version of the IBM Tivoli Monitoring probe rules files and ObjectServer triggers are also copied to the system by the installation program and will replace the previous versions. The `tivoli_eif.rules` file will also be deleted in the omnibus directory under the event synchronization installation directory because it is no longer used.

If you are using a Netcool/OMNIBus multitier architecture, upgrade the Tivoli Monitoring Event Synchronization component with each ObjectServer in the aggregation tier. If you are using a single-tier architecture, upgrade the event synchronization component with each of your ObjectServers.

Upgrading from a wizard

Follow these steps to upgrade the event synchronization component via the installation wizard.

About this task

The tasks in this section can be completed by either the Netcool/OMNIBus Administrator or the IBM Tivoli Monitoring Administrator.

Procedure

1. On the host where you installed the previous version of the event synchronization component (normally the host of the Netcool/OMNIBus ObjectServer), start the event synchronization upgrade installation.
 - **Windows** Double-click the ESUpgrade30Win32.exe file in the tec subdirectory on the IBM Tivoli Monitoring Tools DVD or DVD image.
 - **Linux** **UNIX** Change to the tec subdirectory on the IBM Tivoli Monitoring Tools DVD or DVD image, and run the following command:
`ESUpgrade30operating_system.bin`

where *operating_system* is the operating system you are installing on (aix, HP11, Linux, linux390, or Solaris).

For example, run the following command on an AIX computer:
`ESUpgrade30Aix.bin`
2. Click **Next** on the Welcome window.
3. Select **I accept the terms in the license agreement**, and click **Next**. A window is displayed indicating the installation is about to proceed.
4. Click **Next** to proceed with the installation. A progress indicator shows the progress of the installation.
5. Click **Finish** to exit the installer.

Upgrading existing IBM Tivoli Monitoring events to use new OMNIBus attributes

If you are upgrading from IBM Tivoli Monitoring Version 6.2.2 (or earlier) to the latest version of IBM Tivoli Monitoring, you must upgrade the existing monitoring events in the Netcool/OMNIBus ObjectServer alerts.status table to use a new OMNIBus attribute.

Before you begin

The task in this section should be completed by the Netcool/OMNIBus Administrator.

About this task

To upgrade you apply the itm_migrate.sql file provided with the IBM Tivoli Monitoring Event Synchronization component to the ObjectServer.

The command to configure OMNIBus pipes the SQL command set into the SQL command-line tool and performs the updates to the ObjectServer.

Procedure

Update the Object Server database with the following commands:

- **Windows**

```
%OMNIHOME%\..\bin\redist\isql -U username  
-P password  
-S server_name  
< path_to_file\itm_migrate.sql
```

Where:

\$OMNIHOME

Is the system-defined variable defining the installation location of OMNIBus.

username

Is the OMNIBus Object Server user name.

password

Is the OMNIBus Object Server password.

server_name

Is the OMNIBus Object Server name defined for process control.

path_to_file

Is the fully qualified path to the specified SQL file.

- **UNIX**

```
$OMNIHOME/bin/nco_sql -user username  
-password password  
-server server_name  
< path_to_file/itm_migrate.sql
```

Where:

\$OMNIHOME

Is the system-defined variable defining the installation location of OMNIBus.

username

Is the OMNIBus Object Server user name.

password

Is the OMNIBus Object Server password.

server_name

Is the OMNIBus Object Server name defined for process control.

path_to_file

Is the fully qualified path to the specified SQL file. The SQL files are located in the event_sync_install_dir/omnibus directory where event_sync_install_dir is the location where the IBM Tivoli Monitoring Event Synchronization component is installed.

Replacing the default deduplication trigger

With previous versions of the triggers provided by the IBM Tivoli Monitoring event synchronization component, the default deduplication trigger was overwritten in the Netcool/OMNIBus ObjectServer. The overwriting causes you to lose any customizations you made to the trigger. The latest version of the event synchronization component provides a deduplication trigger specific to Tivoli Monitoring called itm_deduplication.

About this task

If your default deduplication trigger has IBM Tivoli Monitoring logic, follow these steps to apply the default deduplication trigger to the Netcool/OMNIBus ObjectServer. The tasks in this section should be completed by the Netcool/OMNIBus Administrator.

Procedure

1. Open the automation.sql file located at %OMNIHOME%\etc\automation.sql (Windows) or \$OMNIHOME/etc/automation.sql (UNIX).
2. Locate the command that creates the deduplication trigger. For example:

```
create or replace trigger deduplication
group default_triggers
priority 1
comment 'Deduplication processing for ALERTS.STATUS'
before reinsert on alerts.status
for each row
begin
    set old.Tally = old.Tally + 1;
    set old.LastOccurrence = new.LastOccurrence;
    set old.StateChange = getdate();
    set old.InternalLast = getdate();
    set old.Summary = new.Summary;
    set old.AlertKey = new.AlertKey;
    if (( old.Severity = 0) and (new.Severity > 0))
    then
        set old.Severity = new.Severity;
    end if;
end;
go
```
3. Copy this command to a temporary file. For this example, the temporary file is /tmp/dedup.sql.
4. Edit the command in the temporary file to add the line in bold type so that the deduplication trigger ignores IBM Tivoli Monitoring events and does not change the summary and severity of those events:

```
for each row
when (new.Type != 20) and (new.Type != 21)
begin
```
5. Run the command to replace the deduplication trigger.

- **Windows**

```
%OMNIHOME%\..\bin\redist\isql -U username -P password -S server_name < C:\tmp\dedup.sql
```

- **UNIX**

```
$OMNIHOME/bin/nco_sql -user username -password password -server server_name < /tmp/dedup.sql
```

where:

OMNIHOME

Is the system-defined variable defining the installation location of OMNIBus

username

Is the OMNIBus Object Server user name

password

Is the OMNIBus Object Server password

server_name

Is the OMNIBus Object Server name defined for process control.

Updating the EIF Probe Rules

The Netcool/OMNIBus Probe for Tivoli EIF should be updated with the latest versions of the IBM Tivoli Monitoring probe rules files. This task should be performed by the Netcool/OMNIBus administrator.

IBM Tivoli Monitoring requires the Netcool/OMNIBus Probe for Tivoli EIF version 10 or later. This release of the EIF probe provides a `tivoli_eif.rules` file that includes rules files for other products. You should use the probe's version of `tivoli_eif.rules` instead of the `tivoli_eif.rules` file that was provided by IBM Tivoli Monitoring releases prior to V6.2.3. The latest version of the IBM Tivoli Monitoring probe rules are in the `itm_event.rules` include file that is included by the probe's version of `tivoli_eif.rules`.

If you are using a `tivoli_eif.rules` file from a previous release of IBM Tivoli Monitoring and you customized the rules, determine what customizations were made to the `tivoli_eif.rules` file that you backed up before upgrading the EIF Probe. Then perform one of the following actions:

- Merge your customized rules for IBM Tivoli Monitoring events into the `itm_event.rules` provided with the latest version of IBM Tivoli Monitoring. Or:
- Determine if you can use the `itm_event.rules` as it is, and put your IBM Tivoli Monitoring event customization in the new `itm_custom_override.rules` file that is included by `itm_event.rules` to make upgrades easier in the future.

To update the EIF probe to use the latest IBM Tivoli Monitoring rules, see “Configuring the Netcool/OMNIBus EIF probe” on page 760.

Customizing Event Integration

After verifying that event integration is working as expected, additional configuration tasks are available to customize the event integration between IBM Tivoli Monitoring and Netcool/OMNIBus.

Table 139. Additional configuration tasks after event integration installation

Task	Architecture type	Administrator
“Configuring IBM Tivoli Monitoring Situation Update Forwarder event flow to OMNIBus” on page 773.	Bidirectional	Netcool/OMNIBus
“Changing the configuration values of the IBM Tivoli Monitoring Situation Update Forwarder” on page 773.	Bidirectional	Netcool/OMNIBus
“Updating the IBM Tivoli Monitoring Situation Forwarder to forward event status updates to additional monitoring servers” on page 774.	Bidirectional	Netcool/OMNIBus
“Customizing how the IBM Tivoli Monitoring OMNIBus triggers handle event status updates from the monitoring servers” on page 775.	Bidirectional	Netcool/OMNIBus
“Changing the default acknowledgment timeout used when sampled events are deleted or cleared in Netcool/OMNIBus” on page 778.	Bidirectional	Netcool/OMNIBus
“Editing the configuration parameters that the hub monitoring server uses to forward events” on page 778.	Unidirectional and bidirectional	IBM Tivoli Monitoring
“Specifying which situations forward events to Netcool/OMNIBus” on page 780.	Unidirectional and bidirectional	IBM Tivoli Monitoring Operator
“Understanding and Customizing the Event Contents” on page 787.	Unidirectional and bidirectional	IBM Tivoli Monitoring

Table 139. Additional configuration tasks after event integration installation (continued)

Task	Architecture type	Administrator
“Converting from unidirectional architecture to bidirectional architecture” on page 781.	Bidirectional	Netcool/OMNIBus
“Customizing event status processing behavior when agent switching is used or the agent goes offline” on page 782.	Unidirectional and bidirectional	IBM Tivoli Monitoring
“Creating a Netcool/OMNIBus WebGUI tool to launch from WebGUI to the Tivoli Enterprise Portal” on page 785.	Unidirectional and bidirectional	Netcool/OMNIBus

Configuring IBM Tivoli Monitoring Situation Update Forwarder event flow to OMNIBus

This section outlines the scenarios and process for configuring IBM Tivoli Monitoring Situation Update Forwarder event flow to OMNIBus

If you are using the bidirectional architecture, the IBM Tivoli Monitoring Situation Update Forwarder can send events to Netcool/OMNIBus in the following cases:

- The Situation Update Forwarder sends events to Netcool/OMNIBus if it detect errors when it is processing event status updates initiated from the OMNIBus Object Server.
- The Situation Update Forwarder sends an event to Netcool/OMNIBus to update the default acknowledgement timeout when the `sitconf.sh` or `sitconf.cmd` commands are executed.

To send events from the Situation Update Forwarder to Netcool/OMNIBus, update the values for the following parameters in the `event_sync_install_dir/omnibus/errorevent.conf` file:

ServerName
ServerPort

where:

event_sync_install_dir

Is the location where the IBM Tivoli Monitoring Event Synchronization Component is installed on your Netcool/OMNIBus ObjectServer system.

ServerName

Is the name of the computer where the EIF probe is running.

ServerPort

Is the listening port for the EIF probe. The default value is 9998.

Changing the configuration values of the IBM Tivoli Monitoring Situation Update Forwarder

After installation of IBM Tivoli Event Synchronization, you can change the configuration values of the Situation Update Forwarder that were specified during installation.

See Table 134 on page 741 and Table 135 on page 742 for the list of configuration values that can be modified. Use the procedure in this section to change the configuration parameters of the IBM Tivoli Monitoring Situation Update Forwarder when you are using the bidirectional architecture.

If you want to change any of the settings for the IBM Tivoli Monitoring Situation Update Forwarder on the Netcool/OMNIBus ObjectServer system, use the `sitconfig.sh` command on Linux/UNIX or the `sitconf.cmd` command on Windows. You have two options for running this command:

- Manually modify the IBM Tivoli Monitoring Situation Update Forwarder configuration file for event synchronization (named `situpdate.conf` by default and located in `<event_sync_install_dir>/etc`) and then run: `sitconfig.sh update <config_filename>`.
- Run the `sitconfig.sh` command directly, specifying only those settings that you want to change.

See the for Tivoli Netcool/OMNIBus commands for the full syntax of this command.

After you change the configuration of the event synchronization, you must manually stop and restart the IBM Tivoli Monitoring Situation Update Forwarder process from the `<event_sync_installdir>/bin` directory. On a Windows operating system, use the **stopSUF.cmd** and **startSUF.cmd** commands. On operating systems such as UNIX, use the **stopSUF.sh** and **startSUF.sh** commands.

Updating the IBM Tivoli Monitoring Situation Forwarder to forward event status updates to additional monitoring servers

If you are using the bidirectional architecture and have multiple hub monitoring servers forwarding situation events to the same Netcool/OMNIBus ObjectServer, the IBM Tivoli Monitoring Situation Update Forwarder must be defined with the list of monitoring servers that it should send event status updates to.

When you install the IBM Tivoli Monitoring Event Synchronization Component and specify the IBM Tivoli Monitoring Situation Forwarder configuration parameters, you specify the SOAP URL of one hub monitoring server.

To add additional hub monitoring servers to the list that can receive event status updates from the Netcool/OMNIBus ObjectServer, issue the `sitconfuser` command as described in the following steps. Ensure the IBM Tivoli Monitoring Situation Update Forwarder is configured for both the fully qualified host name and the short host name of the hub monitoring server.

If you have configured the Hot Standby feature, you must configure the host name information for both the primary and secondary hub monitoring servers.

This task is performed on the system where the Netcool/OMNIBus ObjectServer and IBM Tivoli Monitoring Situation Update Forwarder are installed.

- **Windows** Change to the `<event_sync_installdir>\bin` directory and enter the following command:
`sitconfuser.cmd add serverid=server userid=user password=password
pathc=path_to_conf_file type=OMNIBUS`

where:

<event_sync_installdir>

Is the directory where the IBM Tivoli Monitoring event synchronization component is installed on your Netcool/OMNIBus ObjectServer system.

server

Is the fully qualified host name of the hub monitoring server.

user

Is the user ID used to access the computer where the hub monitoring server is running.

password

Is the password of the user ID specified by the user parameter.

path_to_conf_file

Is the directory containing the `situser.conf` file. By default, the `situser.conf` file is in the `<event_sync_installdir>\etc` directory.

Repeat this command to add short host name information for the same hub monitoring server by specifying the short host name value for the `serverid` parameter.

- **UNIX** Change to the `<event_sync_installdir>/bin` directory and enter the following command:
`sitconfuser.sh add serverid=server userid=user password=password
pathc=path_to_conf_file type=OMNIBUS`

where:

<event_sync_installdir>

Is the directory where the IBM Tivoli Monitoring event synchronization component is installed on your Netcool/OMNIBus ObjectServer system.

server

Is the fully qualified host name of the monitoring server.

user

Is the user ID used to access the computer where the monitoring server is running.

password

Is the password of the user ID specified by the user parameter.

path_to_conf_file

Is the directory containing the `situser.conf` file. By default, the `situser.conf` file is in the `<event_sync_installdir>/etc` directory.

Repeat this command to add short host name information for the same monitoring server by specifying the short host name value for the `serverid` parameter.

Repeat this step for each monitoring server that you want to add.

You can also delete monitoring servers, see the for the full syntax of this command.

After you change the configuration of the event synchronization, you must manually stop and restart the IBM Tivoli Monitoring Situation Update Forwarder process from the `<event_sync_installdir>/bin` directory. On Windows operating systems use `stopSUF.cmd` and `startSUF.cmd`. On UNIX operating systems use `stopSUF.sh` and `startSUF.sh`.

Customizing how the IBM Tivoli Monitoring OMNIBus triggers handle event status updates from the monitoring servers

You can configure how the IBM Tivoli Monitoring OMNIBus triggers handle acknowledgment expiration and resurface status update events from the hub monitoring servers when the bidirectional architecture is used.

You can also configure how stop situation event status updates are handled for pure situation events for both the unidirectional and bidirectional architecture. For more information about handling event status updates, see “Event behavior” on page 700.

It might also be necessary to edit the procedure that sends event status updates from Netcool/OMNIBus to IBM Tivoli Monitoring if Netcool/OMNIBus is installed into a directory path containing spaces or the Situation Update Forwarder will be run by a user other than root.

Configuring acknowledgement expiration and resurface status update event behavior

This section outlines the process for configuring acknowledgement expiration and resurface status update event behavior.

The **get_config_parms** procedure in the `<event_sync_install_dir>/omnibus/itm_proc.sql` file defines three configuration parameters:

```
set sit_ack_expired_def_action = 'REJECT'  
set sit_resurface_def_action = 'ACCEPT'  
set situpdate_conf_file = 'situpdate.conf'
```

Note: When unidirectional architecture is used, the `itm_proc.sql` file must be configured to accept the acknowledgment expiration and resurface status update events from the hub monitoring servers.

The `sit_ack_expired_def_action` variable defines the action to be taken for an event by the OMNIBus server when an acknowledgment expiration status update is received for an event from a hub monitoring server. An acknowledgment expiration status update occurs if the event was acknowledged in the Tivoli Enterprise Portal with a timeout and the event is still true when the timeout expires. The default action is to reject the request if the event is still in the acknowledged state in the Netcool/OMNIBus ObjectServer. In this case, OMNIBus sends information back to the hub monitoring server to change the state of the event from Acknowledgement Expired to Acknowledged. If you want to change the action taken by the OMNIBus server to Accept the acknowledgment expiration, modify the statement in `itm_proc.sql` to set `sit_ack_expired_def_action = 'ACCEPT'`. In this case, the event is deacknowledged in Netcool/OMNIBus but still has the acknowledged with expiration status in the hub monitoring server.

The variable `sit_resurface_def_action` defines the action to be taken by the OMNIBus server when a situation event has resurfaced in IBM Tivoli Monitoring. An event is resurfaced when it has been deacknowledged in the Tivoli Enterprise Portal Situation Event Console. The default action of the OMNIBus server is to Accept this request and Deacknowledge the event. If you would like to change the action taken by OMNIBus server to Reject the resurface of the event, modify the statement to set `sit_resurface_def_action = 'REJECT'` in `itm_proc.sql`. If the action is REJECT, OMNIBus then sends information back to the hub monitoring server to change the state of the event back to Acknowledged and the event remains in the acknowledged state in the Netcool/OMNIBus ObjectServer.

The variable `situpdate_conf_file` specifies the name of the configuration file to be used by the IBM Tivoli Monitoring Situation Update Forwarder. If you would like to change the name of the configuration file, modify the statement to set `situpdate_conf_file = 'newname.conf'`.

After modifying `itm_proc.sql`, issue the following command on your Netcool/OMNIBus ObjectServer:

- **Windows**

```
%OMNIHOME%\..\bin\redist\isql -U <username>  
-P <password>  
-S <server_name>  
< <path_to_file>\itm_proc.sql
```
- **Linux** **UNIX**

```
$OMNIHOME/bin/nco_sql -user <username>  
-password <password>  
-server <server_name>  
< <path_to_file>/itm_proc.sql
```

Where:

OMNIHOME

Is the system-defined variable defining the installation location of OMNIBus.

username

Is the OMNIBus ObjectServer user name.

password

Is the OMNIBus ObjectServer password.

server_name

Is the OMNIBus ObjectServer name defined for process control.

path_to_file

Is the fully qualified path to the specified SQL file.

Configuring stop situation event behavior for pure situations

When the IBM Tivoli Monitoring triggers in Netcool/OMNIBus process a situation stop event, the triggers clear all events for the situation that are detected by the remote monitoring server specified by the ITMThruNode OMNIBus attribute. However, you can configure the IBM Tivoli Monitoring triggers to ignore situation stop events for pure events.

To configure this behavior, you must update the following statement in the `itm_event_clear` trigger from:

```
set puresitstop = 'CLOSE';
```

to:

```
set puresitstop = 'OPEN';
```

You can either use the Netcool/OMNIBus Administrator to edit the `itm_event_clear` trigger or modify the trigger in the `itm_db_update.sql` file provided with the IBM Tivoli Monitoring event synchronization component and then re-load the `itm_db_update.sql` file into Netcool/OMNIBus ObjectServer. After modifying `itm_db_update.sql`, issue the following command on your Netcool/OMNIBus ObjectServer:

- **Windows**

```
%OMNIHOME%\..\bin\redist\isql -U <username>  
-P <password>  
-S <server_name>  
< <path_to_file>\itm_db_update.sql
```
- **Linux** **UNIX**

```
$OMNIHOME/bin/nco_sql -user <username>  
-password <password>  
-server <server_name>  
< <path_to_file>/itm_db_update.sql
```

Where:

OMNIHOME

Is the system-defined variable defining the installation location of OMNIBus.

username

Is the OMNIBus ObjectServer user name.

password

Is the OMNIBus ObjectServer password.

server_name

Is the OMNIBus ObjectServer name defined for process control.

path_to_file

Is the fully qualified path to the specified SQL file.

Editing the eventcmd procedure

This section lists possible modifications to the `eventcmd` procedure.

The `eventcmd` procedure in the `<event_sync_install_dir>/omnibus/itm_proc.sql` file might also need the following modifications:

- The value of the executable variable must not contain any spaces. If the Situation Update Forwarder was installed in a directory with spaces, change the executable variable setting from `$EVENTCMD` to a path for the `eventcmd.bat` file that does not contain spaces, for example, on Windows:
`C:\Progra~1\IBM\SitForwarder\omnibus\eventcmd.bat.`
- Change the host variable to reflect the actual host name on which the Object Server is running.
- You might have to change the user and group variables from the default settings if the executable cannot be run as root.

Editing the writeitmcmd procedure

This section lists possible modifications to the writeitmcmd procedure.

You might also need to make the following modifications to the writeitmcmd procedure in the <event_sync_install_dir>/omnibus/itm_proc.sql file:

- Change the host variable to reflect the actual host name on which the Object Server is running.
- You might have to change the user and group variables from the default settings if the executable cannot be run as root.

Changing the default acknowledgment timeout used when sampled events are deleted or cleared in Netcool/OMNIBus

When a situation event from a sampled situation is forwarded to the Tivoli Netcool/OMNIBus ObjectServer and that event is later deleted or cleared in the ObjectServer, the behavior of the bidirectional event synchronization architecture is to send a request to the hub Tivoli Enterprise Monitoring Server to acknowledge the situation with a specified timeout.

The reason for this behavior is that you cannot close sampled situation events unless the monitoring agent determines the situation condition is no longer true.

If the acknowledgment timeout of the situation expires and the situation is still true, IBM Tivoli Monitoring sends an acknowledgement expiration status update event to Netcool/OMNIBus. The status update event causes a new event to be opened in the Netcool/OMNIBus ObjectServer and the Netcool/OMNIBus operator is notified that the event condition has not been resolved. If the situation condition becomes false, then the event is closed in IBM Tivoli Monitoring, and the event remains cleared in the Netcool/OMNIBus ObjectServer.

The default acknowledgment expire time for sampled situations is 59 minutes. This default time can be changed in the situation timeouts configuration file on the ObjectServer (sit_timeouts.conf). Also, expiration times for individual situations can be configured in this file. After editing this file, you can have the expire times dynamically loaded into the ObjectServer by using the sitconf.sh refresh (UNIX) or sitconf.cmd refresh (Windows) command in <event_sync_installdir>/bin.

Note: You must perform the task in the section “Configuring IBM Tivoli Monitoring Situation Update Forwarder event flow to OMNIBus” on page 773 before running the sitconf command.

Editing the configuration parameters that the hub monitoring server uses to forward events

Edit the hub monitoring server Tivoli Event Integration Facility EIF configuration file to customize the configuration parameters that specify up to five failover Netcool/OMNIBus EIF probes or to adjust size of the event cache.

About this task

When the Tivoli Event Integration Facility (EIF) has been enabled on the hub monitoring server and the default EIF server (Tivoli Enterprise Console Event Server or Netcool/OMNIBus EIF probe) and port numbers have been specified, the EIF configuration file is updated with the information. This is the default EIF receiver of forwarded situation events.

See “Configuring the hub monitoring server to forward events” on page 764 for instructions on configuring the monitoring server to enable the Tivoli Event Integration Facility.

Take the following steps to edit the EIF configuration file:

Procedure

1. Open the `om_tec.config` file:
 - a. In the Manage Tivoli Monitoring Services window, right-click Tivoli Enterprise Monitoring Server and click **Advanced → Edit EIF Configuration**.
 - b. Open `Install_dir/tables/tems_name/TECLIB/om_tec.config` in a text editor where `Install_dir` is the directory where IBM Tivoli Monitoring is installed and `tems_name` is the name of the hub monitoring server supplied during monitoring server installation.
2. Edit any of the event server configuration parameters for the event integration facility.
3. When you are finished editing `om_tec.config`, save the file.
4. You must restart the hub monitoring server or, alternatively, you can use the `tacmd refreshTECinfo` command to complete the updates without having to restart the monitoring server. To use this command, log in to the command-line interface with `tacmd login`, then run `tacmd refreshTECinfo -t eif` to complete the EIF configuration. For more information on this command, see the *IBM Tivoli Monitoring Command Reference*.

Table 140. Supported Netcool/OMNIBus Probe for Tivoli EIF configuration parameters for the Event Integration Facility (EIF)

Parameters	Value	Remarks
ServerLocation=	tec_server_addr	Host name or IP address of the OMNIBus EIF probe. To provide event failover, you can indicate up to five default EIF probes, separating each with a comma. When the default EIF probe is unavailable, the situation event goes to the next server in the list.
ServerPort=	[port:0]	OMNIBus EIF probe listening port, which is 9998 by default. Specify 0 if the OMNIBus EIF probe uses the port mapper. If you specified multiple server locations, add the corresponding port numbers here, separating each with a comma.
EventMaxSize=	4096	Maximum number of characters allowed in the event. This parameter is disabled by default. To enable it, remove the pound symbol (#) at the beginning of the line.
RetryInterval=	5	The number of times to retry connection with the EIF probe before returning an error.
getport_total_timeout_usec=	50500	The number of seconds to continue attempting to connect to the EIF probe port before timing out. The default is 14 hours.
NO_UTF8_CONVERSION=	YES	Events are already in UTF8, no conversion is needed. Must be set to YES.
ConnectionMode=	co	The connection mode.
BufferEvents=	YES	Whether the EIF buffers the event. This must be set to YES.

Table 140. Supported Netcool/OMNIBus Probe for Tivoli EIF configuration parameters for the Event Integration Facility (EIF) (continued)

Parameters	Value	Remarks
BufEvtMaxSize=	4096	Maximum size of the event cache. The default is initially 4096 KB and you can change it here.
BufEvtPath=	./TECLIB/om_tec.cache	Path of the event cache file. The default is ./TECLIB/om_tec.cache.
FilterMode=	OUT	Enable event filtering. This is set to OUT by default.
Filter:	Class=ITM_Generic; master_reset_flag="";	To filter out specific classes, use this keyword. By default, situation events of the class ITM_Generic and those that send no master reset flag are not forwarded.

Specifying which situations forward events to Netcool/OMNIBus

By default, when the hub Tivoli Enterprise Monitoring Server has been configured for the Tivoli Event Integration Facility, events for situations enabled for event forwarding are sent to the default EIF receiver.

Before you begin

The tasks in this section can be completed by the IBM Tivoli Monitoring Tivoli Enterprise Portal Operator.

About this task

See the “Configuring the hub monitoring server to forward events” on page 764 section for details on how to specify the default EIF receiver. If you want to forward situation events to multiple Netcool/OMNIBus Probes for Tivoli EIF, you must use the `tacmd createEventDest` command to create additional event destinations and then use the Tivoli Enterprise Portal Situation Editor to specify which event destinations to use for the situation’s events. For details on the `tacmd createEventDest` command, see the *IBM Tivoli Monitoring Command Reference*.

Also by default, event forwarding is not enabled for a new situation unless you base the new situation definition on an existing situation that already has event forwarding enabled, or you explicitly enable event forwarding for the situation. You can also use the Tivoli Enterprise Portal Situation Editor to ensure that Event Integration Facility forwarding is enabled for each situation whose events should be forwarded to Netcool/OMNIBus.

Complete these steps to specify the destination EIF receiver and severity for a forwarded event:

Procedure

1. In the Tivoli Enterprise Portal Navigator view, either right-click the Navigator item that the situation is associated with and click **Situations** or click **Situation Editor** in the main toolbar.
2. Select the situation whose events should be forwarded to Netcool/OMNIBus.
3. Click the **EIF** tab.
4. Select **Forward Events to an EIF Receiver** to specify that an EIF event is sent for each event that opens for this situation.
5. Select the **EIF Severity** to apply to forwarded events for this situation. <Default EIF Severity> uses the same severity as is used for the situation at this Navigator item.
6. Assign any other EIF receivers:

- To add a destination, select it from the **Available EIF Receivers** list and move to the Assigned list. (After selecting the first destination, you can use Ctrl+click to select other destinations or Shift+click to select all destinations between the first selection and this one.)
 - To remove a destination, select it from the **Assigned EIF Receivers** list and move to the Available list. The **Available EIF Receivers** list shows all of the defined EIF destinations that were created with the `tacmd createEventDest` command.
7. Save the situation definition with your changes by clicking **Apply**, to keep the Situation editor open, or **OK**, to close the Situation editor.

Converting from unidirectional architecture to bidirectional architecture

This section outlines the steps for converting from unidirectional to bidirectional architecture.

About this task

If you are using the unidirectional architecture for event integration between your hub monitoring servers and Netcool/OMNIBus and you decide to switch to the bidirectional architecture, you must perform the following tasks for each of your Netcool/OMNIBus ObjectServers that will be sending event status updates back to a hub monitoring server:

Procedure

1. Ensure you have the IBM Tivoli Monitoring Event Synchronization Component installed on the machine with your Netcool/OMNIBus ObjectServer.
2. Configure the OMNIBus server to support execution from scripts for running the Situation Update Forwarder. For more information, see “Configuring the OMNIBus server for program execution from scripts” on page 763.
3. Define each monitoring server that is forwarding events to Netcool/OMNIBus ObjectServer to the IBM Tivoli Monitoring Situation Update Forwarder and start the IBM Tivoli Monitoring Situation Update Forwarder. For more information, see “Updating the IBM Tivoli Monitoring Situation Forwarder to forward event status updates to additional monitoring servers” on page 774.
4. Update the Netcool/OMNIBus ObjectServer database schema to add the IBM Tivoli Monitoring triggers that send bi-directional event status updates to your hub monitoring servers. Follow the procedure in “Updating the OMNIBus database schema on single-tier or aggregation tier ObjectServers” on page 753 to load the `itm_sync.sql` file. You do not need to load the `itm_db_update.sql` or `itm_proc.sql` files because they were loaded into the ObjectServer when you enabled unidirectional event integration.
5. Edit the `get_config_parms` procedure in Netcool/OMNIBus Object to choose the default bidirectional behavior when an Acknowledgement Expired event status update is sent from the hub monitoring server to Netcool/OMNIBus:
 - a. Use the `nco_config` utility to start Netcool/OMNIBus Administrator.
 - b. Connect to the Netcool/OMNIBus ObjectServer where you are making the change.
 - c. Select **Automation** → **Procedures** from the menu.
 - d. Edit the `get_config_parms` procedure and change:


```
set sit_ack_expired_def_action = 'ACCEPT';
```

to

```
set sit_ack_expired_def_action = 'REJECT';
```
 - e. Save the trigger and exit from Netcool/OMNIBus Administrator.

Customizing event status processing behavior when agent switching is used or the agent goes offline

This section contains the variables for customizing the behavior of event status processing when agent switching is used.

The variables in Table 141 on page 783 can be added to the monitoring server's environment file to customize the behavior of event status processing when agent switching is used or when the agent goes offline. The first two variables help ensure that events are not closed by the agent's primary monitoring server after the agent has switched to its secondary monitoring server. The monitoring server's environment file can be found in these locations:

- **Windows**

ITM_HOME\cms\KBBENV
For example: C:\IBM\ITM\cms\KBBENV

- **Linux** **UNIX**

ITM_HOME/config/tems_hostname_ms_tems_name.config
For example: /opt/IBM/ITM/config/edinburg_ms_labtems.config

For Linux/UNIX systems, you must add the variables to the .config and the .ini files. The name and location of the .ini file is ITM_HOME/config/ms.ini.

- **z/OS**

&shilev.&rtename.RKANPARU(KDSENV)
For example: ITM.SYP1.RKANPARU(KDSENV)

Note: The &shilev and &rtename are variables that correspond to high level qualifiers of the RKANPARU(KDSENV) partitioned dataset. These variables can take 1 to 8 characters.

You must recycle the Tivoli Enterprise Monitoring Server after modifying the environment file for your changes to be picked up.

Table 141. Variables to customize the behavior of event status processing when agent switching is used

Variable	Architecture Type	Details	Administrator
IRA_MIN_NO_DATA_WAIT_TIME	Unidirectional and bidirectional	<p>The minimum time to wait before the monitoring server closes a situation event. This parameter is defined in number of seconds. The default value is zero.</p> <p>By default, after an agent is disconnected from a Tivoli Enterprise Monitoring Server, situations already open will remain open for three situation polling intervals. For example, take two sampled situations, S1 and S2, with intervals of 30 seconds and 15 minutes respectively. Both situations are open when the agent loses connection. Situation S1 closes after at least one minute and 30 seconds. Situation S2 closes after at least 45 minutes. With agent switching, if a situation closes too soon it might generate duplicate events because the agent did not have sufficient time to connect to the backup monitoring server before the primary server closes the original event. This is particularly true for situations with very short polling intervals.</p> <p>In such a scenario you can use the IRA_MIN_NO_DATA_WAIT_TIME variable to set the minimum wait time before a situation is closed. Using the example above, if IRA_MIN_NO_DATA_WAIT_TIME is set to 300 (5 minutes), S1 will close after 5 minutes not 90 seconds. S2 is unaffected and will close after 45 minutes as before.</p> <p>Note:</p> <ol style="list-style-type: none"> 1. You should set this variable in the environment file for all of your monitoring servers. 2. IRA_MIN_NO_DATA_WAIT_TIME does not prevent a situation from closing if an agent switches from a remote monitoring server that has been stopped. The variable only applies when an agent switches from a running remote monitoring server to the backup remote monitoring server because of a connectivity issue or when the agent is stopped. 	IBM Tivoli Monitoring

Table 141. Variables to customize the behavior of event status processing when agent switching is used (continued)

Variable	Architecture Type	Details	Administrator
CMS_SIT_TIME_VALIDATION	Unidirectional and bidirectional	<p>Valid entries are Y or N. The default is N.</p> <p>By default, the monitoring server handles situation events on a first-come-first-serve basis. In a scenario where agent switching is enabled an agent might send events through two different monitoring servers. The events that arrive first might not necessarily be the earlier events if one of the monitoring servers encountered connection issues. This generally has little impact on situation event processing, except when a monitoring server is shutdown and some situations might be closed prematurely even though the agent is already connected to a different monitoring server.</p> <p>You must perform two actions to avoid this scenario:</p> <ol style="list-style-type: none"> 1. All monitoring server hosts should synchronize time, preferable through Internet Time Protocol (ITP) clients. 2. You should add the CMS_SIT_TIME_VALIDATION=Y variable to all monitoring server environment files. This switches the LCLTMSTMP column in situation events to use UTC time instead of local time, which is then used to determine event order. 	IBM Tivoli Monitoring

Table 141. Variables to customize the behavior of event status processing when agent switching is used (continued)

Variable	Architecture Type	Details	Administrator
CMS_SIT_CHECK_NODESTS	Unidirectional and bidirectional	<p>Valid entries are Y or N. The default is N.</p> <p>This variable is only applicable for users of event integration with Tivoli Enterprise Console and Omnibus. When the CMS_SIT_CHECK_NODESTS variable is set to Y in the environment file, the hub monitoring servers check the agent status whenever a close status update event is forwarded to Netcool/Omnibus. The CMS_SIT_CHECK_NODESTS variable should only be added to the hub monitoring server environment file. If the agent is offline the close status update event is tagged with a special OFFLINE indicator in the situation_eventdata EIF slot.</p> <p>If you do not want events to be closed in Netcool/OMNIBus when an agent goes offline, you can customize the EIF probe rules to ignore close events where the situation_status slot is set to N and the situation_eventdata EIF slot is set to OFFLINE. See “Customizing the rules file” on page 762 for details on how to add customizations to the EIF probe rules. You should also consider setting the IRA_MIN_NO_DATA_WAIT_TIME environment variables described in this table so that close status events are not sent to Netcool/OMNIBus until after the agent offline condition has been detected. (If the default agent heartbeat interval is used, it can take between 10 to 20 minutes before the monitoring server detects that the agent is no longer online.)</p> <p>Note: You should only set the CMS_SIT_CHECK_NODESTS variable in the environment file of your hub monitoring server.</p>	IBM Tivoli Monitoring

Creating a Netcool/OMNIBus WebGUI tool to launch from WebGUI to the Tivoli Enterprise Portal

If you want your Netcool/OMNIBus WebGUI users to launch the Tivoli Enterprise Portal to see additional details about events shown in the WebGUI, you can create a WebGUI tool that launches the Tivoli Enterprise Portal from the Netcool/OMNIBus Event List UI.

The Tivoli Enterprise Portal will be launched in the context of the managed system associated with the event. This support allows WebGUI users to quickly determine and resolve issues in their IT environment.

Follow the instructions in the IBM Tivoli Netcool/OMNIBus Information Center to create event management tools for the WebGUI and choose Script as the tool type. For more details on the tool creation procedure, see the Netcool/OMNIBus information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm.

For the script command, use the format example below that supports your environment.

Note: The examples below use http and port number 1920. If you use https or a different port number to access the Tivoli Enterprise Portal Server, you must customize the script commands below to reflect that environment.

Scenario 1: Single hub monitoring server is sending events to the Netcool/OMNIBus ObjectServer. The Tivoli Enterprise Portal Server and hub monitoring server are on the same host. Use the following script command format for this scenario:

```
var str_appllabel = '{@ITMAppLabel}';
if (!(str_appllabel == "A:P:S") && !(str_appllabel == "A:P:L" )
{ window.open("http://{@ITMHostname}:1920//cnp/kdh/lib/cnp.html?managed_system_name={@ITMSitOrigin}"); }
```

Scenario 2: Single hub monitoring server is sending events to the Netcool/OMNIBus ObjectServer. Tivoli Enterprise Portal Server and the hub monitoring server are on different hosts. Use the following script command format for this scenario:

```
var str_appllabel = '{@ITMAppLabel}';
if (!(str_appllabel == "A:P:S") && !(str_appllabel == "A:P:L" )
{ window.open("http://<teps-hostname>:
1920//cnp/kdh/lib/cnp.html?managed_system_name={@ITMSitOrigin}"); }
```

In this example, replace <teps-hostname> with the IP address or hostname of your Tivoli Enterprise Portal Server.

Scenario 3: Multiple hub monitoring servers are sending events to the same Netcool/OMNIBus ObjectServer. Use the following script command format for this scenario:

```
var str_appllabel = '{@ITMAppLabel}';
if ( ! (str_appllabel == "A:P:S") && !(str_appllabel == "A:P:L" )
{

var strtems={@ITMHostname};
if( strtems == "<tems_host1>" )
{ window.open("http://<teps_host1>:1920//cnp/kdh/lib/cnp.html?managed_system_name={@ITMSitOrigin}"); }
else if( strtems == "<tems_host2>")
{ window.open("http://<teps_host2>:1920//cnp/kdh/lib/cnp.html?managed_system_name={@ITMSitOrigin}"); }
}
```

Change the following in this script command example:

- <tems_host1> to the IP address or the host name of the server where the first hub monitoring server exists.
- <teps_host1> to the IP address or the host name of the server where the first portal server exists.
- <tems_host2> to the IP address or the host name of the server where the second hub monitoring server exists.
- <teps_host2> to the IP address or the host name of the server where the second portal server exists.

If you have more than two hub monitoring servers sending events to the Netcool/OMNIBus ObjectServer, add additional *else if* checks for each hub monitoring server.

You can configure the Netcool/OMNIBus WebGUI and Tivoli Enterprise Portal to authenticate users from an external LDAP repository. If you do this configuration and add the Single Sign-On (SSO) capability between both servers (WebGUI and Tivoli Enterprise Portal server), you can open and work on both consoles with the same user credentials (authenticated by LDAP server). You do not need to login separately to each console.

In order to configure Tivoli Enterprise Portal console for Single Sign-On, see the *IBM Tivoli Monitoring Administrator's Guide*. To configure the WebGUI console for Single Sign-On, see the Netcool/OMNIBus Information Center http://publib.boulder.ibm.com/infocenter/tivihelp/v8r1/topic/com.ibm.tivoli.namomnibus.doc/welcome_ob.htm.

Understanding and Customizing the Event Contents

From the **Situation editor** EIF tab of the Tivoli Enterprise Portal, you can create map definitions for situation events sent to the EIF Probe. The EIF Slot Customization window, which is opened from the EIF tab, is used to customize how situation events are mapped to forwarded EIF events, thus overriding the default mapping between situation event attributes and EIF attribute slots.

With IBM Tivoli Monitoring you can also use the `itm_custom_override.rules` file to customize the mapping of EIF event attributes to OMNIBus attributes. If you want to map an EIF event attribute to a new OMNIBus attribute, you must add the attribute to the ObjectServer alerts.status table before updating the `itm_custom_override.rules` file.

Note: Neither EIF Slot Customization or custom rules should change the values of the attributes that are used by the IBM Tivoli Monitoring triggers in the Netcool/OMNIBus ObjectServer. For details on attributes that should not be modified, see “Default mapping of situation events to OMNIBus events.”

Table 142. Common tasks for understanding and customizing the contents of an event

Task	Architecture Type	Administrator
See “Default mapping of situation events to OMNIBus events” and “Generic mapping for agent-specific slots” on page 792 for details on the default mappings of event attributes to EIF slots and OMNIBus attributes.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIBus
See the <i>IBM Tivoli Monitoring Tivoli Enterprise Portal User’s Guide</i> or the Tivoli Enterprise Portal online help for details on using the EIF slot customization function to customize how situation event attributes are mapped to EIF attribute slots.	Unidirectional and bidirectional	IBM Tivoli Monitoring
“Adding event classes for new and updated agents to the MCS Attribute service used by the EIF Slot Customization function” on page 793.	Unidirectional and bidirectional	IBM Tivoli Monitoring
“About customizing the format of the event message attribute” on page 795.	Unidirectional and bidirectional	IBM Tivoli Monitoring
“Localizing the event message attribute” on page 796.	Unidirectional and bidirectional	IBM Tivoli Monitoring
See “Customizing the rules file” on page 762 for details on how to customize the mapping of IBM Tivoli Monitoring EIF event attributes to Netcool/OMNIBus ObjectServer attributes.	Unidirectional and bidirectional	IBM Tivoli Monitoring and Netcool/OMNIBus

Default mapping of situation events to OMNIBus events

You can use the mapping information in this topic when you forward a situation event to the Tivoli Netcool/OMNIBus ObjectServer and want to write probe rules or SQL procedures and triggers in the ObjectServer.

The situation event forwarder generates an event integration facility (EIF) event with an event class based on the attribute group used in the situation. When the situation event is forwarded to the EIF Probe, the Tivoli Netcool/OMNIBus EIF probe translates the EIF event into an OMNIBus alert format. The EIF event contains all of the base attributes described by the *Omegamon_Base* class and its parent class EVENT. The EIF event also contains the agent specific attributes from the attribute group referenced by the situation definition when the event is opened. The agent specific attributes are also called extended slots in the Tivoli Enterprise Portal EIF Slot Customization function. For more information on agent-specific slots included in events, see “Generic mapping for agent-specific slots” on page 792.

Note: Agent-specific attributes are not included in any event status updates (including the event status update message that is sent when the event is closed in IBM Tivoli Monitoring).

Omegamon_Base is described as follows:

```
Omegamon_Base ISA EVENT
DEFINES {
cms_hostname: STRING;
cms_port: STRING;
integration_type: STRING;
master_reset_flag: STRING;
appl_label: STRING;
situation_name: STRING;
situation_origin: STRING;
situation_displayitem: STRING;
situation_time: STRING;
situation_status: STRING;
situation_eventdata: STRING;
situation_type: STRING;
situation_thrunode: STRING;
situation_group: STRING;
situation_fullname: STRING; }; END;
```

As part of the generic mapping for these situations, the IBM Tivoli Monitoring event forwarder assigns associated values for each of the *Omegamon_Base* event class attributes when forwarding an event to the EIF Probe. In addition to these event class attributes, values are assigned to the host name, origin, severity, and message attributes that are inherited from the EVENT class as shown in Table 143.

Table 143. Base IBM Tivoli Monitoring EIF attribute slots.

EIF slot attributes	Values and meaning
appl_label	Indicates if the event was sent from a monitoring agent. For events sent from a monitoring agent, this attribute is set to A:P:S if the event is a private situation event or is set to A:P:L if the event is a lifecycle event. For events sent from a hub monitoring server, a value is not provided for this slot.
cms_hostname	TCP/IP host name of the hub Tivoli Enterprise Monitoring Server that forwards the event.
cms_port	The hub monitoring server port on which the SOAP web service is listening.
ClassName	ITM_ + the name of the attribute group that the situation definition is based on, for example: ITM_NT_Process.
fqhostname	The fully qualified host name, if available, of the managed system where the event originated. ¹
hostname	The TCP/IP host name of the managed system where the event originated, if available. ¹

Table 143. Base IBM Tivoli Monitoring EIF attribute slots. (continued)

EIF slot attributes	Values and meaning
integration_type	Indicator to help OMNIbus performance. <ul style="list-style-type: none"> • N for a new event, the first time the event is raised • U for update event, subsequent event status changes The integration_type value is used solely by the OMNIbus synchronization rule to improve its performance. It has no other meaning related with the event.
master_reset_flag	Master reset indicator set for master reset events that indicate a hub monitoring server has been started. Value is NULL for all other events: <ul style="list-style-type: none"> • R for monitoring server recycle master_reset • S for hotstandby master_reset
msg	The situation name and formula when a situation event is opened and the situation name for all other event messages, for example acknowledgement, close, and so on.
origin	The TCP/IP address of the managed system where the event originates, if available. The address is in dotted decimal format. ¹
severity	The EIF severity of the event.
situation_displayitem	Display item of associated situation, if available.
situation_eventdata	Event data attributes in key-value pair format if multiple rows of the attribute group match the situation condition, if the situation does not have a display item, or if multiple rows have the same display item value. The data from the first row that matches the situation condition is not included in this EIF slot. The event data can be truncated because the event integration facility imposes a 2 KB size limit.
situation_group	One or more situation group names (up to 5) that the situation is a member of.
situation_fullname	Display name of the associated situation if it is different from the situation identifier in the situation_name slot.
situation_name	Unique identifier given to the situation.
situation_origin	Managed system name where the situation event originated. It has the same value as sub_source.
situation_status	Status of the situation event. The following values are valid: <ul style="list-style-type: none"> • Y if the situation condition is true and event has been opened in the hub monitoring server. • N if the situation condition has become false and the event is closed in the hub monitoring server. • A if the event has been acknowledged. • E if the event has been resurfaced (deacknowledged). • F if the acknowledgment timeout for the event has expired. • P if the situation has been stopped. • D if the situation has been deleted. • X the situation includes an error.

Table 143. Base IBM Tivoli Monitoring EIF attribute slots. (continued)

EIF slot attributes	Values and meaning
situation_time	Timestamp of the situation event.
situation_type	Indicator of whether the IBM Tivoli Monitoring situation that caused the event is a sampled or pure situation.
situation_thrunode	The hub or remote Tivoli Enterprise Monitoring Server that the managed system is connected to.
source	ITM
sub_source	The origin managed system name for the associated situation.
Note: 1. If there is an IBM Tivoli Monitoring firewall gateway between the managed system and its monitoring server, these slots contain the hostname and IP address of the firewall gateway instead of the managed system.	

The IBM Tivoli Monitoring rules for the Tivoli Netcool/OMNIBus EIF probe map the EIF attribute slots of the situation event into ObjectServer attributes, which are defined in the alerts.status table of the Netcool/OMNIBus ObjectServer. See Table 144 for details. Although the EIF event also contains the agent specific attribute group slots when the event is opened, these slots are not mapped to ObjectServer attributes by the probe rules.

Table 144. Mapping of EIF attribute slots to OMNIBus attributes

EIF slot attribute	OMNIBus attribute
situation_name + situation_origin + situation_displayitem + event_class	Identifier (for ITMProblem)
situation_name + situation_origin + situation_displayitem + event_class + ITMResolution	Identifier (for ITMResolution)
situation_name	AlertKey
situation_origin	ITMSitOrigin
situation_origin	Node
situation_origin	NodeAlias
source	Agent
default	Type (20) (for ITMProblem)
situation_status = "P" and integration_type = "U"	Type (21) (for ITMResolution)
situation_status = "D" and integration_type = "U"	Type (21) (for ITMResolution)
situation_status = "N" and integration_type = "U"	Type (21) (for ITMResolution)
situation_displayitem	ITMDisplayItem
situation_status	ITMStatus
situation_thrunode	ITMThruNode
situation_time	ITMTime
situation_type	ITMSitType
situation_eventdata	ITMEventData
cms_hostname	ITMHostname
master_reset_flag	ITMResetFlag
integration_type	ITMIntType
ClassName	AlertGroup

Table 144. Mapping of EIF attribute slots to OMNibus attributes (continued)

EIF slot attribute	OMNibus attribute
msg	Summary
"tivoli_eif probe on "+hostname()	Manager
6601 if integrating with Tivoli Business Service Manager 87722 if the event was sent by the hub monitoring server 87723 if the event was sent from an IBM Tivoli Monitoring agent	Class
severity FATAL / 60 = Critical CRITICAL / 50 = Critical MINOR / 40 = Minor WARNING / 30 = Warning UNKNOWN / 10 = Indeterminate	Severity
getdate	LastOccurrence/FirstOccurrence
date	TECDate
repeat_count	TECRepeatCount
fqhostname	TECFQHostname
hostname	TECHostname
cms_port	ITMPort
situation_fullname	ITMSitFullName
situation_group	ITMSitGroup
appl_label	ITMApplLabel

Note: If you use the virtualization or predictive rules files provided with Netcool/OMNibus, those rules might set some of the OMNibus attributes differently than the settings in the table above, for example, Node, NodeAlias, Class, and Summary.

You should not customize the EIF probe rules or use the EIF Slot Customization window of the Situation Editor to change any of the situation EIF attribute slots that are used by the IBM Tivoli Monitoring triggers in the Netcool/OMNibus ObjectServer to process IBM Tivoli Monitoring events. The EIF slots and OMNibus attributes that should not be customized are shown in table Table 145.

Table 145. EIF attribute slots to OMNibus attributes, not to be customized

EIF Slot	OMNibus attribute
appl_label	ITMApplLabel
cms_hostname	ITMHostname
cms_port	ITMPort
integration_type	ITMIntType
master_reset_flag	ITMResetFlag
situation_displayitem	ITMDisplayItem
situation_eventdata	ITMEventData
situation_name	AlertKey
situation_origin	ITMSitOrigin
situation_status	ITMStatus
situation_thruNode	ITMThruNode

Note: The situation_origin EIF slot is also mapped to the Node and NodeAlias OMNIBus attributes. However, since the IBM Tivoli Monitoring triggers do not use the Node and NodeAlias attributes, you can customize their values in the itm_custom_override.rules file. However, if you are using the virtualization or predictive rules files provided with Netcool/OMNIBus, verify that your customizations of Node and NodeAlias are consistent with the values set by these rules.

Generic mapping for agent-specific slots

Generic mapping is used for events for which there is no event mapping file and is based on the target event class that is identified by the situation definition.

For situation events that do not have a custom event mapping specified for the forwarded event, an event is generated with a unique ClassName value based on the attribute group used in the situation. The ClassName attribute of the EIF event is set to a combination of ITM_ plus the attribute group name associated with the situation. For example, a situation using the NT_Process attribute group generates an EIF event with ClassName set to ITM_NT_Process. The EIF event contains the base attribute slots described in Table 143 on page 788.

Additional agent-specific event slot values are populated with attribute values from the situation event data. The EIF slot names are the attribute names from the attribute group. For example, a situation using the Process_CPU attribute causes generation of a slot called process_cpu in the EIF event forwarded to the OMNIBus EIF probe. If an agent-specific attribute name conflicts with a slot name in the Tivoli Enterprise Console EVENT class or Omegamon_Base class, then the agent product code associated with the attribute group, for example: knt_, is prepended to the attribute name to form a unique slot name.

By default, the agent specific EIF event slots are not mapped to OMNIBus attributes. However, you can perform custom mapping of these EIF slots to OMNIBus attributes in the itm_custom_override.rules file. You can either map the agent specific EIF slots to name-value pairs in the ExtendedAttr OMNIBus attribute, or map them to a new OMNIBus attribute. Also some agents, for example, ITCAM for Transactions agents, might provide their own rules and .sql file for mapping agent specific attributes to new or existing OMNIBus attributes. (If you create a new OMNIBus attribute, you must add the attribute to the ObjectServer alerts.status table before updating the rules file.) However, it is important to note that the agent specific slot values are only included in the EIF event that is sent when the event is opened and not in any subsequent status update events, for example when the event is acknowledged or closed. If you are using bidirectional event synchronization between the monitoring server and Netcool/OMNIBus and you map agent specific slots to a new OMNIBus attribute, you must also update the IBM Tivoli Monitoring table and triggers that are used to cache sampled events that are cleared or deleted in the OMNIBus UI. The event data is cached so that it can be used to fully populate a new event if the event is re-opened when the acknowledge timeout expires in the monitoring server. See Table 125 on page 700 for more information about the behavior of events originating from a hub monitoring server.

Customize the list of OMNIBus attributes that are cached

Perform the procedure in this section to customize the list of OMNIBus attributes that are cached when a sampled event is cleared or deleted after the new attribute has been added to the ObjectServer alerts.status table.

Procedure

1. Open the itm_event_cache.sql file using a text editor. This file is available in the event_sync_install_dir/omnibus where event_sync_install_dir is the directory where the IBM Tivoli Monitoring event synchronization component is installed.
2. Find "--NewAttributeGoesHere type(size)". Copy the line, insert a new line immediately below, and paste the copied line. In the newly added line delete the characters "--", replace "NewAttributeGoesHere type(size)" with the new OMNIBus attribute name, type, and size.

- Find "--statusr.NewAttributeGoesHere,". Copy the line, insert a new line immediately below, and paste the copied line. In the newly added line delete the characters "--", replace "NewAttributeGoesHere" with the appropriate OMNIBus attribute name. Do not delete the comma (,) at the end of the line.
- Find "--set new.NewAttributeGoesHere = cachedevent.NewAttributeGoesHere ;". Copy the line, insert a new line immediately below, and paste the copied line. In the newly added line delete the characters "--", replace "NewAttributeGoesHere" with the appropriate OMNIBus attribute name. Do not delete the semicolon (;) at the end of the line.
- Execute the itm_event_cache.sql file to update the OMNIBus database schema by using the instructions in "Updating the OMNIBus database schema on single-tier or aggregation tier ObjectServers" on page 753. For complex situations, the situation definition can involve more than one attribute group. In this case, the EIF event class used is derived from the first attribute group encountered in the situation event data of the triggering situation. For example, if a situation is written for the NT_Process and Local Time attribute groups, NT_Process being the first attribute group, the EIF event class ITM_NT_Process is used. Additional event slots are generated based on the attributes of the first attribute group only.

Table 146. Special characters for attribute groups and names in EIF events generated from forwarded situation events.

Character:	Converts to:
<uppercase> (applies only to attribute name)	<lowercase> (applies only to attribute name)
% percent sign	pct_
I/O	io
/ forward slash	_per_
\ backward slash	_ (underscore)
<space>	_ (underscore)
(open parenthesis	_ (underscore)
) close parenthesis	
< open pointed bracket	_ (underscore)
> close pointed bracket	

All strings and time stamp types are mapped to STRING types, and all integer types are mapped to INTEGER in the event class definition. No default values are assigned to the attribute slots. Attributes that have a specified non-zero scale/precision value are mapped to the string type of REAL.

Note: If you are mapping from an attribute to a slot and the resulting slot name has a trailing underscore, the trailing underscore is removed in the final slot name. Final slot names never have a trailing underscore.

Adding event classes for new and updated agents to the MCS Attribute service used by the EIF Slot Customization function

The EIF Slot Customization facility uses the MCS Attribute Service to present a list of predefined event classes in the Event class name list of the EIF Slot Customization window. The Event class name list of the EIF Slot Customization window is available through the EIF tab of the Tivoli Enterprise Portal Situation editor.

Only the event classes belonging to the OS agents are predefined and they are in an MCS Attribute Service jar file. When a new type of agent is added into the Tivoli Management Services infrastructure or

a new event class is added to an updated version of an agent, you must generate a new MCS XML file and point the Tivoli Enterprise Portal Server to the new XML file before the new event classes will appear in the Event class name list.

See the *IBM Tivoli Monitoring Tivoli Enterprise Portal User's Guide* or the Tivoli Enterprise Portal online help for details on using the EIF slot customization function.

To generate a new MCS XML file, install the Tivoli Enterprise Console Event Definition Generator (TEDGEN) tool supplied on the IBM Tivoli Monitoring Tools DVD. The TEDGEN tool should be installed on a distributed computer where the Hub monitoring server or portal server is installed. The TECLIB directory of the Hub monitoring server or portal server contains the agent BAROC files that are processed by the TEDGEN tool.

Note: The definitions in MCS XML file supersede those defined in the shipped MCS Attribute Services jar file (they are not merged). To obtain a MCS XML file that contains both the event classes definitions of the OS agents as well as new or updated agents, be sure all the BAROC definitions for the OS agents and other agents used in your environment are all in the TECLIB directory of your portal server or monitoring server before running the TEDGEN utility to generate the MCS XML file.

Installation and Configuration

This section outlines the steps for adding event classes for new and updated agents to the MCS Attribute service used by the EIF Slot Customization function.

Before you begin

1. Install the TEDGEN tool from the IBM Tivoli Monitoring Tools DVD on either the Hub monitoring server or portal server. The tool is located in the `tec/tedgen` directory on the Tools DVD and the installation and configuration instructions are in the `README.txt` file in the same directory.
2. If you installed the TEDGEN tool on a portal server on Linux or UNIX, you must also perform the following configuration steps:
 - a. Create the `Install_dir/tables/cicatrsg/TECLIB` directory if it does not already exist, where `Install_dir` is the directory where IBM Tivoli Monitoring is installed.
 - b. Copy the `om_tec.baroc` and `kib.baroc` files from the `Install_dir/arch/cq/TECLIB` directory to the `Install_dir/tables/cicatrsg/TECLIB` directory where `arch` is the architecture directory for the portal server.

These steps assume that you have installed the TEDGEN utility on the computer where the Hub Tivoli Enterprise Monitoring Server or the Tivoli Enterprise Portal Server are installed and you have installed the latest application support for the agents whose event slots will be customized.

Note:

1. If you are generating the MCS XML file on the portal server on Linux or UNIX, the BAROC files are not present by default and you must install them by running the `install.sh` script on the portal server computer and selecting the **Install TEMS support for remote seeding** option. See "Installing application support files on a computer with no monitoring server" on page 298 for more details. This action places the BAROC files on the portal server under the `Install_dir/tables/cicatrsg/TECLIB` directory.
2. If you installed the TEDGEN utility with a portal server on Windows, verify that the portal server's TECLIB directory contains the `.baroc` files for the agents whose events you want to customize. Not all agents include their `.baroc` files in their Tivoli Enterprise Portal Server application support. If an agent's `.baroc` file is not present, you can copy it from the Hub monitoring server's TECLIB directory.

Perform these steps to run the TEDGEN utility:

Procedure

1. If you are running the TEDGEN utility on Windows, issue the following command:

```
set CANDLE_HOME=Install_dir
```

where *Install_dir* is the directory where IBM Tivoli Monitoring is installed.

2. If you are running the TEDGEN utility on UNIX, issue the following command:

```
export CANDLEHOME=Install_dir
```

where *Install_dir* is the directory where IBM Tivoli Monitoring is installed.

3. Change to the TEDGEN_Install_dir/scripts directory where TEDGEN_Install_dir is the directory where the TEDGEN utility is installed.
4. Run the TEDGEN tool to create a new MCS XML file:

- **Windows**

```
tedgen -itmDir Install_dir\{CMS|CNPS}  
TECLIB -id server_id -xmlPath output_xml_file_path
```
- **Linux** **UNIX**

```
tedgen -itmDir  
Install_dir/tables/{tems_name|cicatrsg}/TECLIB  
-id server_id -xmlPath output_xml_file_path
```

Where:

Install_dir

Is the directory where the monitoring server or portal server is installed.

server_id

Is any string. The EIF slot customization function does not use the value of this string.

output_xml_file_path

Is the name of the MCS XML file that should be created by the tool. Absolute and relative paths are supported.

Example: In the following example, the hub monitoring server named mytems has the BAROC files in the TECLIB directory. The output file goes to the same directory and is named tems.xml.

```
tedgen -itmDir C:\IBM\ITM\CMS\TECLIB -id mytems -xmlPath tems.xml
```

5. If the TEDGEN tool was run on the Hub monitoring server, copy the newly generated MCS XML file to the computer where the Tivoli Enterprise Portal Server is installed.
6. Edit the portal server environment file to specify the path to the XML file:
 - a. On Windows systems: In the Manage Tivoli Monitoring Services window, right-click **Tivoli Enterprise Portal Server** and click **Advanced → Edit ENV File** to open the kfwenv file in the text editor.
 - b. On Linux/UNIX systems: Open *Install_dir/config/cq.ini* in a text editor.
 - c. Locate the KFW_MCS_XML_FILES environment variable and type = (equal sign) followed by the path to the MCS XML file.
 - d. Save and close the environment file.
 - e. On Windows systems, restart the portal server. On Linux/UNIX systems, reconfigure the portal server and then start it.

About customizing the format of the event message attribute

The OMNIBus EIF probe maps the msg attribute slot in the EIF event sent from the Tivoli Enterprise Monitoring Server into the Summary attribute of the ObjectServer. The Summary attribute gives you a descriptive way of looking at an alert in OMNIBus.

The situation name alone does not provide detailed event identification where you have large numbers of like-events from various sources. Therefore, the msg EIF slot is expanded to include the following event attributes when an event is opened:

Situation-Identifier [(*formula*) ON *Managed-System-Name* ON *DISPLAY-ITEM* (*threshold Name-Value pairs*)]

where:

Situation-Name

The unique identifier of the situation.

formula

The formula tells how the situation is evaluated.

Managed-System-Name

The agent or the managed system.

DISPLAY-ITEM

The identifier that triggered the situation if there is more than one instance. This is optional and is used only if a display item is specified in the situation definition.

threshold Name-Value pairs

The raw data that the situation uses to evaluate whether it is triggered.

Examples:

```
NT_Critical_Process [(Process_CPU > 4 AND Thread_Count > 50)
ON IBM-AGX02:NT
(Process_CPU = 8 AND Thread_Count = 56)]
```

```
NT_Disk_Full [(Free_Megabytes < 1000000)
ON "IBM-AGX02:NT"
ON D: (Free_Megabytes = 100)]
```

You can use the EIF Slot Customization window of the Tivoli Enterprise Portal Situation Editor to override the default msg slot setting.

When customizing the msg attribute slot, the Literal value column can be used to define a custom message template. The message template can consist of fixed message text and variable substitution references, or symbols. The symbol can be base or extended slot data, or a special reference to the situation formula. Base slots are those that are included in all forwarded events, such as `situation_name`. Extended slots are those specific to the attribute group used by the situation definition. See the following syntax:

- For an extended slot, use the fully qualified attribute name (**\$Attribute_Table.Attribute_Name\$**).
- For a base slot, use the variable name that is not fully qualified (no . periods) unless it is the situation formula symbol.
- For a situation formula, use **\$formula\$**.

These characters are not supported: less than, greater than, quotation mark, single quotation mark, and ampersand. The Literal value column cannot be used to define a message template if a value is selected in the Mapped attribute column.

See the Tivoli Enterprise Portal online help or the *IBM Tivoli Monitoring Tivoli Enterprise Portal User's Guide* for more information about using the EIF slot customization function.

Localizing the event message attribute

Edit the `KMS_OMTEC_GLOBALIZATION_LOC` variable on the hub monitoring server to enable globalization of the EIF event msg attribute slot that is mapped to the OMNIbus Summary attribute by the OMNIbus EIF probe rules.

About this task

By default, this variable is set to United States English and the msg attribute slot contains United States English messages. Take these steps to edit the variable to enable any language packs that are installed in your hub monitoring server environment:

Procedure

1. On the computer where the hub Tivoli Enterprise Monitoring Server is installed, open the KBBENV file:
 - Start Manage Tivoli Monitoring Services, right-click **Tivoli Enterprise Monitoring Server** and click **Advanced → Edit ENV file**.
 - In a text editor, open the `<install_dir>/config/ <tems_name>_ms_<address>.cfg` file, where `<install_dir>` is the directory where IBM Tivoli Monitoring is installed, `<tems_name>` is the value of the hub monitoring server supplied during the monitoring server configuration, and `<address>` is the IP address or fully qualified name of the hub monitoring server computer.
2. Locate (or add) the KMS_OMTEC_GLOBALIZATION_LOC environment variable and enter the desired language and country code in the format `KMS_OMTEC_GLOBALIZATION_LOC=xx_XX`, where
 - xx is the language.
 - XX is the country code: de_DE, en_US, en_GB, es_ES, fr_FR, it_IT, ja_JP, ko_KR, pt_BR, zh_CN, or zh_TW.

For example, `KMS_OMTEC_GLOBALIZATION_LOC=pt_BR` (for Brazilian Portuguese) or `KMS_OMTEC_GLOBALIZATION_LOC=zh_CN` (for Simplified Chinese).
3. Save and close the monitoring server environment file.
4. Restart the hub monitoring server.

Chapter 31. Setting up event forwarding to Tivoli Enterprise Console

Tivoli Enterprise Console events can be forwarded from Tivoli Monitoring to Tivoli Enterprise Console Version 3.9. The events are forwarded from the hub monitoring server to the Tivoli Enterprise Console server or Tivoli Enterprise Console gateway. Make sure that firewall ports are opened between the hub monitoring server and Tivoli Enterprise Console servers. By default, the Tivoli Enterprise Console uses port 5529.

To view updates to the forwarded situation events in the Tivoli Enterprise Portal, you need to install the event synchronization component on the event server. The event synchronization component enables changes in the event status made on the Tivoli Enterprise Console to be reflected on the Tivoli Enterprise Portal.

The following table provides an overview of the tasks required to configure situation event forwarding and synchronization:

Table 147. Tivoli Enterprise Console event synchronization installation and configuration steps

Task	Where to find information
Plan the deployment of your integration.	"Event integration with Tivoli Enterprise Console" on page 800
Gather information required during the installation and configuration processes.	"Information to gather for event forwarding" on page 126
Install the event synchronization component on your event server.	"Installing event synchronization on your event server" on page 804
Install monitoring agent .baroc files on the event server.	"Installing monitoring agent .baroc files on the event server" on page 819
Configure your monitoring server to forward events to Tivoli Enterprise Console.	"Configuring your monitoring server to forward events" on page 820
Start and stop event synchronization on the event server.	"Starting and stopping the Situation Update Forwarder process" on page 821
Modify the configuration of event synchronization as your monitoring needs change.	"Changing the configuration of the event synchronization component on the event server" on page 822
Define additional monitoring servers to the event server.	"Defining additional monitoring servers to the event server" on page 822
Customizing event status processing behavior when agent switching is used or the agent goes offline.	"Customizing event status processing behavior when agent switching is used or the agent goes offline" on page 823
Change the default TCP/IP settings if it is taking a long time to forward events back to the monitoring server.	"Changing the TCP/IP timeout setting on your event server" on page 825

If you are upgrading from a previous IBM Tivoli Monitoring release and already have event synchronization installed for Tivoli Enterprise Console, see "Upgrading to Tivoli Event Synchronization version 3.0.0.0" on page 826.

Event integration with Tivoli Enterprise Console

The scenarios in this section illustrate the various ways to forward situation events from one or more monitoring servers to one or more Tivoli Enterprise Console event servers.

- “One or more hub monitoring servers and a single event server”
- “A single hub monitoring server and multiple event servers” on page 801
- “Multiple hub monitoring servers and multiple event servers in a hub and spoke configuration” on page 802

The Tivoli Enterprise Console product significantly reduces the number of events displayed, enabling you to focus on the most critical, relevant events and manage even the largest, most complex environments. The Tivoli Enterprise Console product gives you the extensive control and flexibility that you need to manage and maintain availability across your enterprise. Managing situation events with the Tivoli Enterprise Console product gives you the following advantages:

- Aggregation of event information from a variety of different sources including those from other Tivoli software applications, Tivoli partner applications, custom applications, network management platforms, and relational database systems
- Pre-configured rules that automatically provide best-practices event management
- Persistence and processing of a high volume of events in an IT environment by:
 - Prioritizing events by their level of importance
 - Filtering redundant or low priority events
 - Correlating events with other events from different sources
 - Root cause analysis and resolution
 - Initiating automatic corrective actions, when appropriate, such as escalation
- Unified system and network management by automatically performing the following event management tasks:
 - Correlating the status of a system or application to the status of the network that it uses
 - Determining if the root cause of a system or application problem is an underlying network failure

Note: If you already have policies that contain emitter activities that send events to the Tivoli Enterprise Console, turning on Tivoli Event Integration event forwarding will result in duplicate events. You can deactivate the emitter activities within policies so you do not have to modify all your policies when you activate Tivoli Event Integration Facility forwarding by using **Disable Workflow Policy/Tivoli Emitter Agent Event Forwarding** when you configure the monitoring server.

Using policies gives you more control over which events are sent and may not want to lose this granularity. Moreover, it is likely the policies that are invoking the Tivoli Enterprise Console emitter are doing little else. If you deactivate these activities, there is no point in running the policy. You may delete policies that are longer required, instead of disabling them.

One or more hub monitoring servers and a single event server

You can configure one or more monitoring servers to forward situation events to an event server.

Figure 155 on page 801 shows multiple hub monitoring servers that are configured to forward situation events to the same event server. The event server sends situation updates based on Tivoli Enterprise Console rules and operator actions back to the hub monitoring server that is associated with that situation.

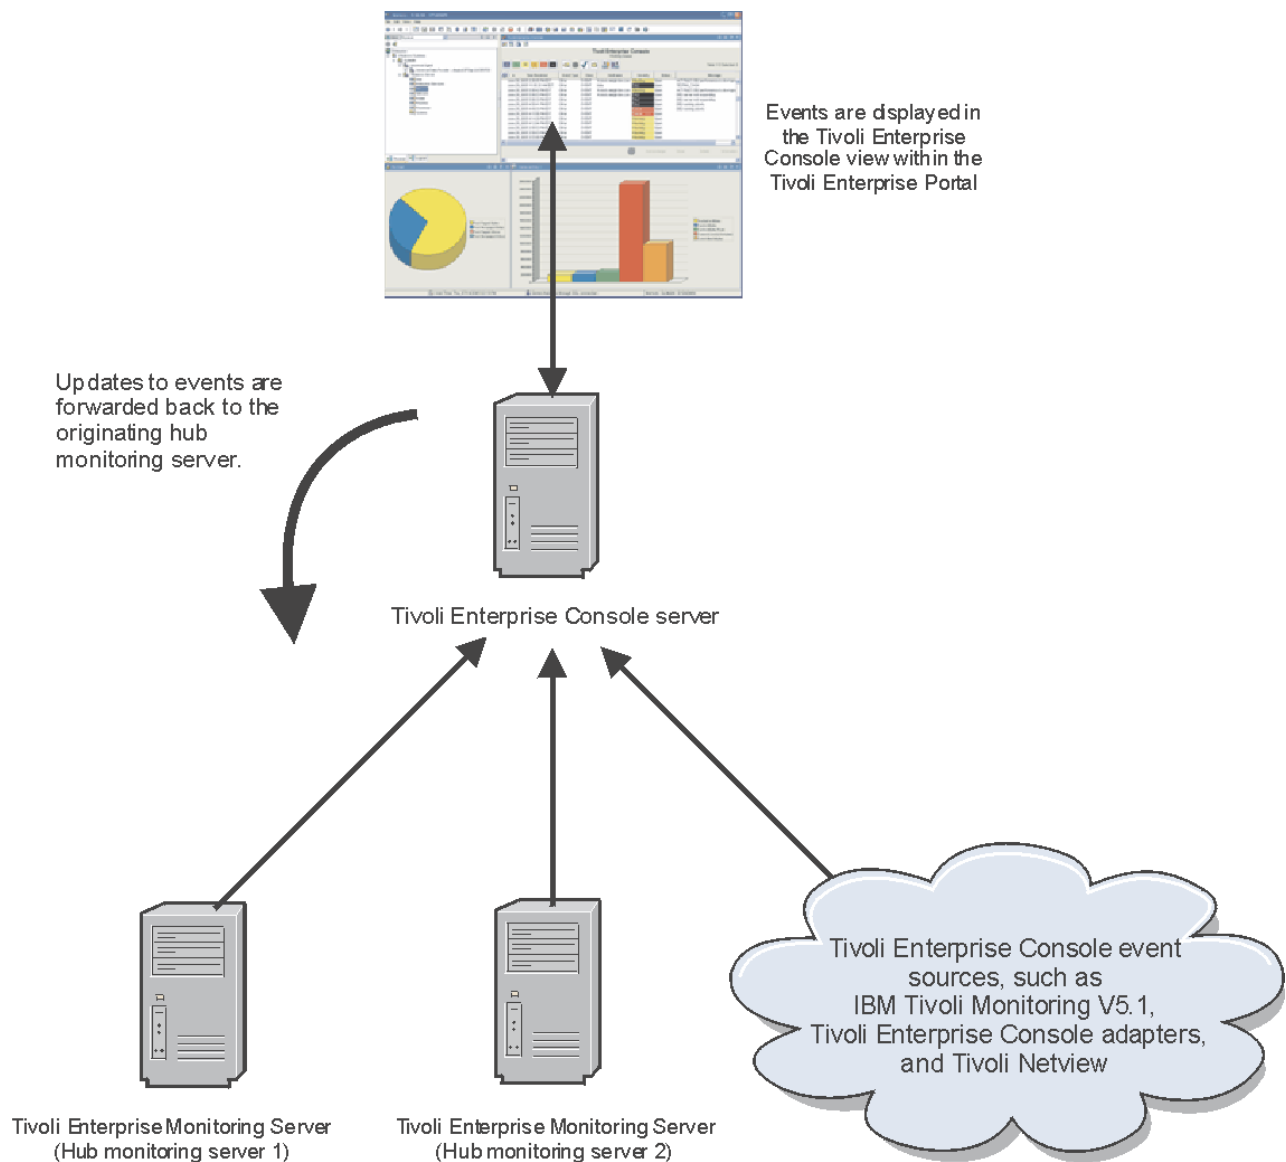


Figure 155. One or more hub monitoring servers connecting to a single event server

A single hub monitoring server and multiple event servers

This section illustrates an environment with a single hub monitoring server and multiple event servers.

Figure 156 on page 802 shows a single hub monitoring server that is configured to forward situation events to multiple event servers. The event servers send situation updates based on Tivoli Enterprise Console rules and operator actions back to the hub monitoring server.

For this configuration, you must install the Tivoli Enterprise Console event synchronization component on each event server, and, for each situation, you must specify the event server to which the situation event is forwarded (see the).

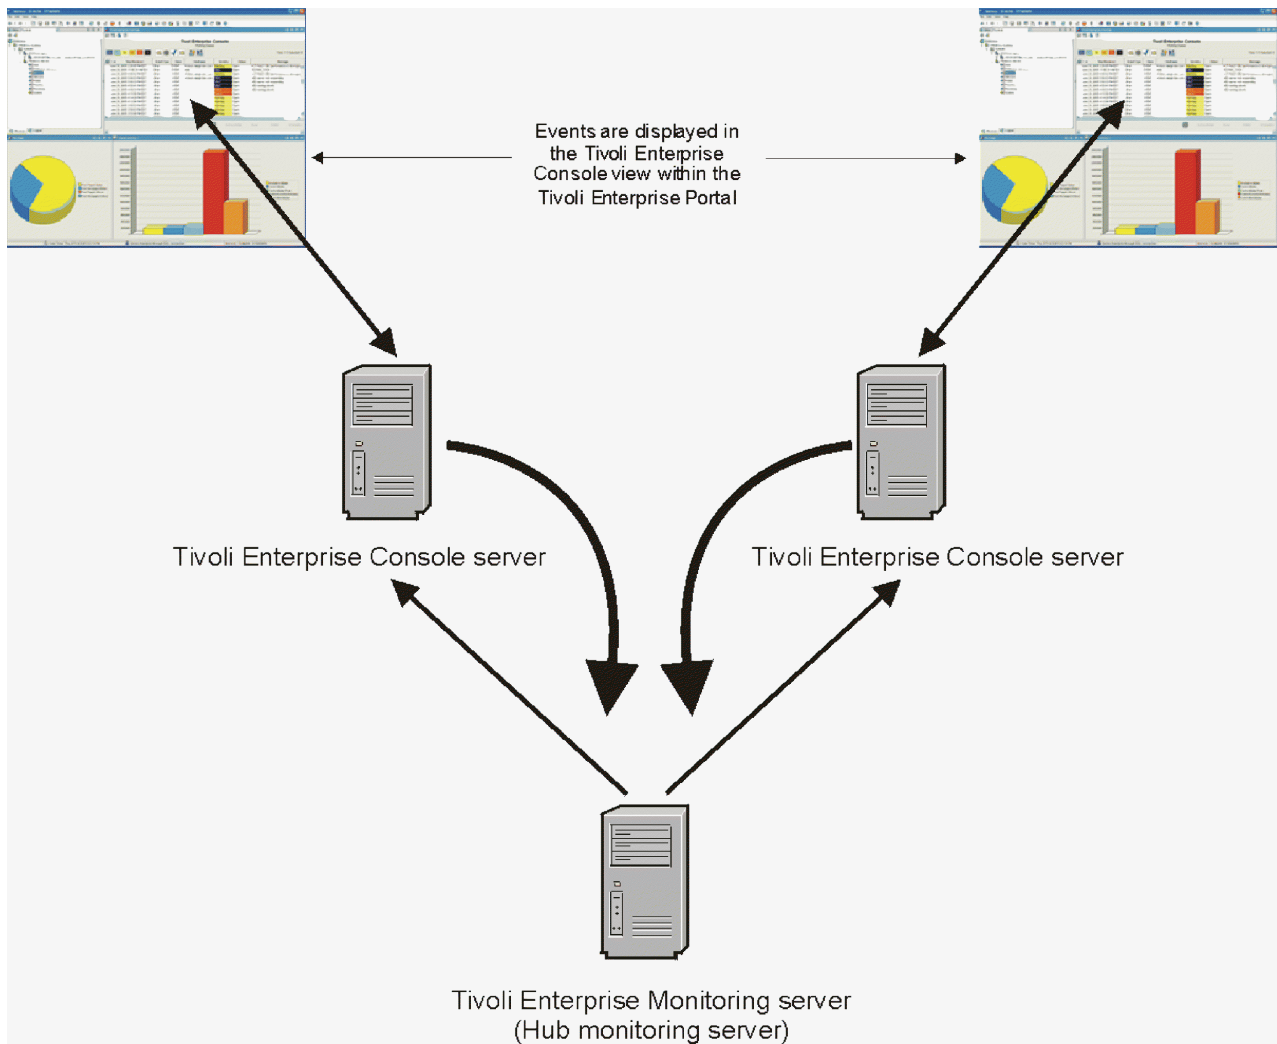


Figure 156. Single hub monitoring server and multiple event servers

Multiple hub monitoring servers and multiple event servers in a hub and spoke configuration

This section illustrates an environment with multiple hub monitoring servers and multiple event servers in a hub and spoke configuration.

Figure 157 on page 803 shows multiple hub monitoring servers that are configured to forward situation events to an event server that is connected to a hub event server. The hub event server sends situation updates based on Tivoli Enterprise Console rules and operator actions back to the hub monitoring server that is associated with that situation.




Figure 157. Multiple hub monitoring servers and multiple event servers in a hub and spoke configuration

Note: This graphic is intended to be an example of one possible scaled configuration for the IBM Tivoli Monitoring and Tivoli Enterprise Console integration. The procedures in this chapter do not provide all of the information needed to set up this sort of configuration.

For this configuration, you must install the Tivoli Enterprise Console event synchronization component on the hub event server. You must also load the `omegamon.baroc` and `Sentry.baroc` files on the spoke event servers, as described in “Modifying an existing rule base” on page 818. In addition, you must load each `.baroc` file for any monitoring agent generating situations that are forwarded to spoke event servers, as described in “Installing monitoring agent `.baroc` files on the event server” on page 819.

Determining when to use the IBM Tivoli Enterprise Console

IBM Tivoli Enterprise Console is a key element of a monitoring environment. Because of its capabilities, the Tivoli Enterprise Console is commonly used in a Manager of Managers role, providing an enterprise level solution for aggregation, consolidation, correlation, and management of events across the enterprise. It is also used as an integration point with which other management applications interact.

Tivoli Enterprise Console provides the following capabilities:

- Aggregation of event information from a large number and variety of different sources including those from other Tivoli software applications, Tivoli partner applications, custom applications, network management platforms, and relational database systems
- Pre-configured rules that automatically provide best-practices event management and root cause determination from an end to end perspective
- Persistence, processing, and access to a high volume of events in an IT environment
- Unified system and network management by automatically performing the following event management tasks:
 - Correlating the status of a system or application to the status of the network that it uses
 - Determining if the root cause of a system or application problem is an underlying network failure

If you are monitoring fewer than 1000 active events and you want to view only situation events (not the other types of events that IBM Tivoli Enterprise Console can monitor), you can use the Situation Event Console in the Tivoli Enterprise Portal. If you are monitoring more than 1000 active events, consider moving to IBM Tivoli Enterprise Console for your event aggregation, and use the Tivoli Enterprise Console view within the Tivoli Enterprise Portal to display the event information. The response time for the Tivoli Enterprise Console view is better than the Situation Event Console view when a large number of events is displayed.

For additional information about the integration with Tivoli Enterprise Console, see “Event synchronization component” on page 13. For additional information about Tivoli Enterprise Console itself, see the Tivoli Enterprise Console information center. For additional information about using the Situation Event Console in the Tivoli Enterprise Portal, see the *IBM Tivoli Monitoring Tivoli Enterprise Portal User's Guide*.

Installing event synchronization on your event server

The installer for the event synchronization component is located in the /tec directory of the IBM Tivoli Monitoring Tools DVD.

There are three methods for installing the component:

- “Installing from a wizard” on page 805
- “Installing from the command-line” on page 809
- “Installing from the command-line using a silent installation” on page 813

When you install the event synchronization component on your event server, the following changes are made:

- If you select to use an existing rule base, the event synchronization .baroc class files (omegamon.baroc and Sentry.baroc [if not present]) and the omegamon.rls rule set file are imported into your existing rule base. If you do not want to modify your existing rule base during the installation, you can choose to manually perform rule base modifications after installation is complete. See “Manually importing the event synchronization class files and rule set” on page 816 for more information.
- For all rule bases that have Sentry.baroc imported, the Sentry2_0_Base class is updated to define additional integration attributes for the situation events received from IBM Tivoli Monitoring.
- A new process, Situation Update Forwarder, is installed along with its supporting binary and configuration files. This process is used to forward updates to the situation events back to the monitoring server. On Windows, a Tivoli Situation Update Forwarder service is also created.

Note:

1. If your IBM Tivoli Enterprise Console event server is running on Windows 2003 and you are planning to install the event synchronization remotely (using a program such as Terminal Services to connect to that Windows 2003 computer), you need to run the **change user /install** command before you run the

installation. This puts the computer into the required "install" mode. After the installation, run the **change user /execute** command to return the computer to its previous mode.

2. If you have a monitoring server on an operating system like UNIX or Linux, you must configure your TCP/IP network services in the `/etc/hosts` file to return the fully qualified host name. See for more information.
3. For a Windows event server, any existing rule base that you use must indicate a relative drive letter (such as `C:\`) as part of its associated path. To verify that your existing rule base contains a relative drive letter, run the following command from a bash environment on your event server:

```
wrb -lsrb -path
```

If the returned path includes something like `hostname:\rulebase_directory`, with no drive letter (such as `C:\`), copy the `ESync3000Win32.exe` file from the `\TEC` subdirectory of the IBM Tivoli Monitoring installation image to the drive where the rule base exists and run the installation from that file.

4. If you are using a Windows event server, if you have any rule base with an associated path that does not contain a relative drive letter and that has the `Sentry2_0_Base` class imported, copy the `ESync3000Win32.exe` file from the `\TEC` subdirectory of the IBM Tivoli Monitoring installation image to the drive where the rule base exists and run the installation from that file.

To verify if you have any rule bases that have an associated path containing no relative drive letter, run the **wrb -lsrb -path** command as described in the previous note.

To determine if your rule bases have the `Sentry2_0_Base` class imported, run the following command against all of your rule bases:

```
wrb -lsrbclass rule_base
```

where *rule_base* is the name of the rule base.

Installing from a wizard

This section outlines the steps for installing event synchronization on your event server from a wizard.

About this task

Use the following steps to install event synchronization using the installation wizard:

Note: The appearance of the window shown in Figure 158 on page 806 indicates that the installer did not find an IBM Tivoli Enterprise Console event server installed on the system, so it expects to install event synchronization for Netcool/OMNIBus. If you see this window when you launch the installation, cancel the installation and install the event server on the system before restarting the installer, or install the synchronization component on another system where an event server is installed.

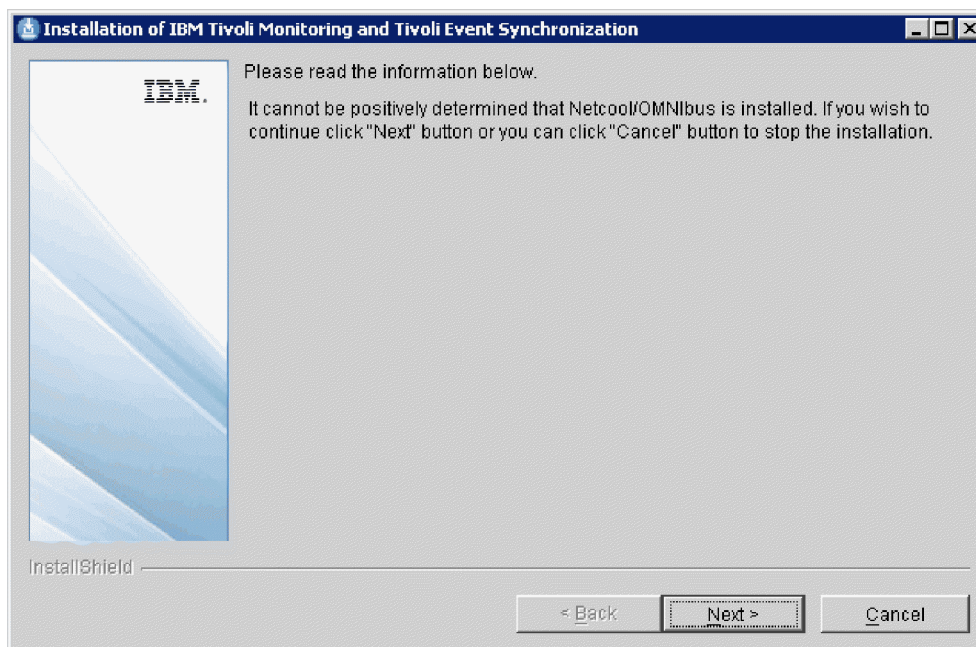


Figure 158. Window shown when no Tivoli Enterprise Console event server is found.

Procedure

1. On the host of the event server, launch the event synchronization installer from the installation media:

Windows Double-click the ESync3000Win32.exe file in the \tec subdirectory on the IBM Tivoli Monitoring Tools DVD or DVD image.

Linux **UNIX** Change to the \tec subdirectory of the IBM Tivoli Monitoring Tools DVD and run the following command:

```
ESync3000operating_system.bin
```

where *operating_system* is the operating system you are installing on (aix, HP11, Linux, linux390, or Solaris). For example, run the following command on an AIX computer:

```
ESync3000Aix.bin
```

2. Click **Next** on the Welcome window.
3. Select **I accept the terms in the license agreement** and click **Next**.
4. Complete the following fields and click **Next**:

Table 148. IBM Tivoli Enterprise Console event synchronization configuration fields

Field	Description
Name of configuration file	The name of the file where event synchronization configuration information is stored. The default name is situpdate.conf.
Number of seconds to sleep when no new situation updates	The polling interval, in seconds. The minimum value is 1, while the default value is 3. If there are no situation events, the Situation Update Forwarder rests for 3 seconds.
Number of bytes to use to save last event	Number of bytes that the long-running process uses when it saves the location of the last event it processes. This value must be an integer. The minimum (and default) value is 50.

Table 148. IBM Tivoli Enterprise Console event synchronization configuration fields (continued)

Field	Description
URL of the CMS SOAP Server	The URL for the SOAP Server configured on the computer where the monitoring server is running. The default value is <code>cms/soap</code> . This value is used to create the URL to which IBM Tivoli Enterprise Console sends event information. For example, <code>http://hostname:port//cms/soap</code> , where <i>hostname</i> is the host name of the monitoring server and <i>port</i> is the port.
Rate for sending SOAP requests to CMS from TEC via web services	The maximum number of event updates sent to the monitoring server at one time. The minimum (and default) value is 10 events.
Level of debug detail for log	The level of information for event synchronization that will be logged. You have the following choices: <ul style="list-style-type: none"> • Low (default) • Medium • Verbose

5. Complete the following information about the files where events will be written and click **Next**:

Table 149. IBM Tivoli Enterprise Console event synchronization configuration fields, continued

Field	Description
Maximum size of any single cache file, in bytes	The maximum permitted size, in bytes, for any one event cache file. The minimum (and default) value is 50000. Do not use commas when specifying this value (specify 50000 instead of 50,000).
Maximum number of caches files	The maximum number of event caches files at any given time. The minimum value is 2, while the default value is 10. When this value is reached, the oldest file is deleted to make room for a new file.
Directory for cache files to be located	The location where event cache files are located. The default locations are as follows: <ul style="list-style-type: none"> • Windows C:\tmp\TME\TEC\OM_TEC\persistence. • UNIX /var/TME/TEC/OM_TEC/persistence

6. Type the following information for each monitoring server with which you want to synchronize events and click **Add**. You must specify information for at least one monitoring server.

Host name

The fully qualified host name for the computer where the monitoring server is running. This name must match the information that will be in events that are issued from this monitoring server.

User ID

The user ID to access the computer where the monitoring server is running.

Password

The password to access the computer.

Confirmation

The same password, for confirmation.

You can add information for up to 10 monitoring servers in this wizard. If you want to add additional monitoring servers, add them after you install them by using the steps provided in “Defining additional monitoring servers to the event server” on page 822.

7. When you have provided information about all of the monitoring servers, click **Next**. You are presented with the options of having the installer automatically perform rule base modifications, or manually performing the modifications after installation is complete (see Table 150).

Table 150. Options for rule base modification

Option	Description
Automatically install rules and classes	The installation wizard will ask for the rule base into which event synchronization class files and rule set will be imported, and automatically execute the rule base commands to do this.
Manually install rules and classes	The installation wizard will not create or update any rule base with event synchronization files. Important: You will have to manually create or update the rule base with event synchronization files after the installation is complete. See “Manually importing the event synchronization class files and rule set” on page 816.

If you select the automatic option, continue with step 8. If you select the manual option, skip to step 11.

8. Specify the rule base that you want to use to synchronize events. You have two choices:

- **Create a new rulebase**
- **Use existing rulebase**

If you select to use an existing rule base, the event synchronization .baroc class files (omegamon.baroc and Sentry.baroc [if not present]) and the omeгамon.rls rule set file are imported into your existing rule base. Also, if Sentry.baroc has already been imported into the existing rule base, the Sentry2_0_Base class is extended to define additional integration attributes for the situation events from IBM Tivoli Monitoring.

- If you are creating a new rule base, type the name for the rule base you want to create and the path to where the new rule base will be located. There is no default location; you must specify a location.
- If you are using an existing rule base, type the name of the rule base.
- If you want to import an existing rule base *into* a new rule base, type the name of the existing rule base in the **Existing rulebase to import** field.

Note: This step is only available if you are creating a new rule base.

9. Click **Next**.
10. If you indicated in the previous step that the installer uses an existing rule base to import the event synchronization class files and rule set, a window is displayed that allows you to specify whether you want the installer to back up the rule base before updating it. If you request a backup, specify both the backup rule base name and backup rule base path. If you leave these fields blank, no backup is made. Click **Next** to proceed to the pre-installation summary panel.
11. Verify the installation location, then click **Next**.
The installation begins.
12. When the installation and configuration steps are finished, a message telling you to stop and restart the event server is displayed.
If you chose to have the installer automatically update the rule base, you are offered the option of having the installer restart the event server for you. Check the box to have the installer restart the server. If you want to restart the event server yourself, leave the box unchecked.
13. Click **OK**.
14. Click **Finish** on the Summary Information window.

Note: If any configuration errors occurred during installation and configuration, you are directed to a log file that contains additional troubleshooting information.

What to do next

Perform the following tasks after the installation is finished:

- Stop and restart the event server for the configuration changes to take effect.
- Install the monitoring agent .baroc files on the event server as described in “Installing monitoring agent .baroc files on the event server” on page 819.
- Configure the monitoring server to forward events to the event server as described in “Configuring your monitoring server to forward events” on page 820.
- If you did not choose to have the rule base updated automatically, update the rule base as described in “Manually importing the event synchronization class files and rule set” on page 816.

Installing from the command-line

This section outlines the steps for installing event synchronization on your event server from the command-line.

About this task

Use the following steps to install the event synchronization from the command-line on your event server:

Procedure

1. Change to the tec directory on the IBM Tivoli Monitoring Tools DVD or DVD image.
2. Run the following command to launch the installation:

Windows

```
ESync3000Win32.exe -console
```

Linux

UNIX

```
ESync3000operating_system.bin -console
```

where *operating_system* is the operating system you are installing on (aix, HP11, Linux, linux390, or Solaris). For example, run the following command on an AIX computer:

```
ESync3000Aix.bin -console
```

The following prompt is displayed:

```
Press 1 for Next, 3 to Cancel or 4 to Redisplay [3]
```

3. Type 1 to start the installation and press Enter.

The following prompt is displayed:

Software Licensing Agreement:

Press Enter to display the license agreement on your screen. Please read the agreement carefully before installing the Program. After reading the agreement, you will be given the opportunity to accept it or decline it. If you choose to decline the agreement, installation will not be completed and you will not be able to use the Program.

4. Press Enter to display the software license agreement.
5. Type 1 and press Enter to accept the license.

The following prompt is displayed:

```
Press 1 for Next, 2 for Previous, 3 to Cancel, or 4 to Redisplay [1]
```

6. Type 1 and press Enter to continue.

The following prompt is displayed:

```
Name of configuration file [situpdate.conf]
```

7. Press Enter to use the default configuration file, situpdate.conf. If you want to use a different configuration file, type the name and press Enter.

The following prompt is displayed:

Number of seconds to sleep when no new situation updates [3]

8. Type the number of seconds that you want to use for the polling interval. The default value is 3, while the minimum value is 1. Press Enter.

The following prompt is displayed:

Number of bytes to use to save last event [50]

9. Type the number of bytes to use to save the last event and press Enter. The default and minimum value is 50.

The following prompt is displayed:

URL of the CMS SOAP server [cms/soap]

10. Type the URL for the monitoring server SOAP server and press Enter. The default value is cms/soap (which you can use if you set up your monitoring server using the defaults for SOAP server configuration).

The following prompt is displayed:

Rate for sending SOAP requests to CMS from TEC via Web Services [10]

11. Supply the maximum number of event updates to send to the monitoring server at one time and press Enter. The default and minimum value is 10.

The following prompt is displayed:

Level of debug for log

[x] 1 low
[] 2 med
[] 3 verbose

To select an item enter its number, or enter 0 when you are finished: [0]

12. Type the level of debugging that you want to use and press Enter. The default value is Low, indicated by an x next to Low.
13. Type 0 when you have finished and press Enter.

The following prompt is displayed:

Press 1 for Next, 2 for Previous, 3 to Cancel, or 4 to Redisplay [1]

14. Type 1 and press Enter to continue.

The following prompt is displayed:

Maximum size of any single cache file, in bytes [50000]

15. Type the maximum size, in bytes, for the cache file and press Enter. The default value is 50000. Do not use commas (,) when specifying this value.

The following prompt is displayed:

Maximum number of cache files [10]

16. Type the maximum number of cache files to have at one time and press Enter. The default value is 10, while the minimum is 2.

On **Windows** the following prompt is displayed:

Directory for cache files to reside [C:/tmp/TME/TEC/OM_TEC/persistence]

On **UNIX** the following prompt is displayed:

Directory for cache files to reside [/var/TME/TEC/OM_TEC/persistence]

17. Type the directory for the cache files and press Enter. The default directory on Windows is C:\tmp\TME\TEC\OM_TEC\persistence; on UNIX, /var/TME/TEC/OM_TEC/persistence.

The following prompt is displayed:

Press 1 for Next, 2 for Previous, 3 to Cancel, or 4 to Redisplay [1]

18. Type 1 and press Enter to continue.

19. The following prompt is displayed:

--- Tivoli Enterprise Monitoring Server 1 ---

Host name []

Type the fully qualified host name for the computer where the monitoring server is running. This name must match the information that is in events issued by this monitoring server. Press Enter.

The following prompt is displayed:

User ID []

20. Type the user ID to use to access the computer where the monitoring server is running and press Enter.

The following prompt is displayed:

Password:

21. Type the password to access the computer and press Enter.

The following prompt is displayed:

Confirmation:

22. Type the password again to confirm it and press Enter.

The following prompt is displayed:

--- Tivoli Enterprise Monitoring Server 2 ---

Host name []

23. Repeat steps 19 on page 810 to 22 for each monitoring server for which you want to receive events on this event server.

When you have provided information for all the monitoring servers *and* you specified information for less than 10 monitoring servers, press Enter to move through the remaining fields defining additional monitoring servers. Do not specify any additional monitoring server information.

The following prompt is displayed:

[X] 1 – Automatically install rules and classes (recommended)
[] 2 – Manually install rules and classes (advanced users)

To select an item enter its number, or 0 when you are finished: [0]

24. If you want to have the installer automatically install the rules and classes, enter 1 and continue with step 25. If you manually install the rules and classes, enter 2 and proceed to step 35.
25. When you see the following prompt, type 1 and press Enter to continue:

Press 1 for Next, 2 for Previous, 3 to cancel or 4 to Redisplay [1]

The following prompt is displayed:

[x] 1 - Create a new rulebase (name and path required)
[] 2 - Use Existing Rulebase (path is optional)

To select an item, enter its number, or press 0 when you are finished: [0]

26. Type 1 to create a new rule base or 2 to use an existing rule base. Press Enter.
27. Type 0 when you are finished and press Enter.
28. If you are creating a new rule base, the following prompt is displayed:

Rulebase Name []

type the name for the rule base and press Enter.

The following prompt is displayed:

Rulebase Path []

29. If you are creating a new rule base, type the path for the new rule base and press Enter.
30. If you are using an existing rule base, the following prompt is displayed:

Rulebase Name []

Type the name of the rule base and press Enter.

31. If you are creating a new rule base, the following prompt is displayed:

Existing rulebase name to import: []

If you want to import an existing rule base into the new rule base, type the name of the existing rule base and press Enter.

The following prompt is displayed:

Press 1 for Next, 2 for Previous, 3 to Cancel, or 4 to Redisplay [1]

32. Type 1 and press Enter to continue.

The following prompt is displayed:

You indicated on the previous dialog that you want to modify an existing rule base. If you want this installer to back up the existing rule base before modifying it please provide a backup rule base name. Leave backup rule base name blank and click Next if you do not want a backup made.

Backup rule base name. []

33. Type the name for the backup rule base and press Enter to continue. If you do not want the installer to back up the existing rule base, press Enter without providing a backup rule base name.

The following prompt is displayed:

If you have provided a backup rule base name you must provide a backup rule base path. NOTE: We append the backup rule base name to the backup rule base path for clarity and easy lookup.

Backup rule base path. []

34. Type the path for the backup rule base and press Enter to continue. If you did not provide a name for a backup rule base, press Enter without providing a rule base path.

The following prompt is displayed:

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

35. Type 1 and press Enter to continue.

The event synchronization is installed.

The following prompt is displayed:

Installation and Configuration has completed. Please stop and restart the Tivoli Enterprise Console Server.

Press 1 for Next, 2 for Previous, 3 to Cancel, or 4 to Redisplay [1]

36. Type 1 and press Enter to continue.

The following prompt is displayed:

Installation and configuration has completed.

Please restart the Tivoli Enterprise Console server for the changes to take effect.

Mark appropriately below to restart the Tivoli Enterprise Console server.

[] 1 - Restart the Tivoli Enterprise Console server to make changes effective

To select an item enter its number, or 0 when you are finished: [0]

The option to automatically restart the Tivoli Enterprise Console is presented only if you chose to have the installer automatically update the rules and classes.

37. If you want the installer to stop and restart the Tivoli Enterprise Console server, type 1 and press Enter. If you want to stop and restart yourself, type 0 and press Enter to continue. The following prompt is displayed:

Press 3 to Finish, or 4 to Redisplay [1]

38. Type 3 to finish and press Enter.

What to do next

Perform the following tasks after the installation is finished:

- If you did not have the installer do it, stop and restart the event server for the configuration changes to take effect.
- Install the monitoring agent .baroc files on the event server as described in “Installing monitoring agent .baroc files on the event server” on page 819.
- Configure the monitoring server to forward events to the event server as described in “Configuring your monitoring server to forward events” on page 820.
- If you did not choose to have the rule base updated automatically, update the rule base as described in “Manually importing the event synchronization class files and rule set” on page 816.

Installing from the command-line using a silent installation

This section outlines the steps for installing event synchronization on your event server from the command-line using a silent installation.

About this task

Use the following steps to install the event synchronization using a silent installation from the command-line on your event server. This installation method runs silently, so you will not see status messages during the actual installation.

Procedure

1. Change to the tec directory of the IBM Tivoli Monitoring Tools DVD or DVD image.
2. Run the following command to generate the configuration file:

Windows

```
ESync3000Win32.exe -options-template filename
```

where *filename* is the name of the configuration file to create, for example, `es_silentinstall.conf`.

UNIX

```
ESync3000operating_system.bin -options-template filename
```

where *operating_system* is the operating system you are installing on (Aix, HP11, Linux, linux390, or Solaris). For example, run the following command on an AIX computer:

```
ESync3000Aix.bin -options-template filename
```

3. Edit the output file to specify the values shown in Table 151 on page 814.

Note:

- a. Remove the pound signs (###) from the beginning of any value that you want to specify.
- b. Do not enclose any values in quotation marks (").
- c. You must specify the following values:

- **configInfoPanel2.fileLocn**
- Information for at least one monitoring server (the **cmdSvrsPnlNotGuiMode.hostname1**, **cmdSvrsPnlNotGuiMode.userID1**, **cmdSvrsPnlNotGuiMode.pswd1**, and **cmdSvrsPnlNotGuiMode.retypePswd1** values)
- **rulebasePanel.chooseNewOrExistingRB**
- **rulebasePanel.rbName**
- **rbInstallTypePanel.rbInstallType**

If you are creating a new rule base, **rulebasePanel.rbPath** is also required. If "manual" is specified for **rbInstallTypePanel.rbInstallType**, all the other **rulebasePanel.*** options are ignored.

If you do not specify any of the other values, the default values are used.

- d. The following values are specified only when installing event synchronization on a Netcool/Omnibus ObjectServer:

```
### -P installLocation
### -W configInfoPanel3.filesize
### -W configInfoPanel3.fileNumber
### -W configInfoPanel3.fileLocn
```

- e. If you specify values, ensure that the value you specify meets the minimum required values. Otherwise, the installation stops and an error is written to the log file.

Table 151. IBM Tivoli Enterprise Console event synchronization configuration values

Value	Description
installLocation	The directory where you want the product to be installed. If the directory path contains spaces, enclose it in double quotes (" "). For example, to install the product to C:\Program Files\My Product, use -P installLocation="C:\Program Files\My Product" Any specified value is ignored if an event server is detected.
configInfoPanel.filename	The name of the file where event synchronization configuration information is stored. The default file name is situpdate.conf.
configInfoPanel.pollingInt	The polling interval, in seconds. The minimum value is 1, while the default value is 3. If there are no situation events, the process that forwards events to IBM Tivoli Enterprise Console rests for 3 seconds.
configInfoPanel.crcByteCnt	Number of bytes that the long running process will use when it saves the location of the last event it processes. This value must be an integer. The minimum (and default) value is 50.
configInfoPanel.cmsSoapURL	The URL for the SOAP Server configured on the computer where the monitoring server is running. The default value is cms/soap. This value is used to create the URL to which IBM Tivoli Enterprise Console sends event information. For example, http://hostname:port//cms/soap, where <i>hostname</i> is the host name of the monitoring server and <i>port</i> is the port.
configInfoPanel.bufFlushRate	The maximum number of event updates sent to the monitoring server at one time. The minimum (and default) value is 10 events.
configInfoPanel.logLevel	The level of debug information for event synchronization that is logged. You have the following choices: <ul style="list-style-type: none"> • Low (default) • Medium • Verbose
configInfoPanel2.filesize	The maximum permitted size, in bytes, for any one event cache file. The minimum (and default) value is 50000. Do not use commas when specifying this value (50,000 instead of 50000).
configInfoPanel2.fileNumber	The maximum number of event caches files at any given time. The minimum value is 2, while the default value is 10. When this value is reached, the oldest file is deleted to make room for a new file.

Table 151. IBM Tivoli Enterprise Console event synchronization configuration values (continued)

Value	Description
configInfoPanel2.fileLocn	The location where event cache files are located. The default locations are as follows: <ul style="list-style-type: none"> Windows C:\tmp\TME\TEC\OM_TEC\persistence. UNIX /var/TME/TEC/OM_TEC/persistence
cmsSvrsPnlNotGuiMode.hostname# Note: The pound sign (#) stands for a number between 1 and 10. For example, "hostname1".	The host name of each monitoring server that will send events to the event server. Specify up to 10 monitoring servers.
cmsSvrsPnlNotGuiMode.userID#	The user ID for the monitoring server, identified in hostname#, to use to access the computer where the monitoring server is running.
cmsSvrsPnlNotGuiMode.pswd#	The password for the user ID used to access the computer where the monitoring server is running.
cmsSvrsPnlNotGuiMode.retypePswd#	The password confirmation for the user ID.
rbInstallTypePanel.rbInstallType	Specifies whether the installer will automatically update a specified rule base, or if a rule base must be manually modified after installation is complete. Specify either automatic or manual. If manual is specified, all other rulebasePanel.* options are ignored.
rulebasePanel.chooseNewOrExistingRB	Specifies whether you are going to create a new rule base or use an existing rule base. Specify either new or existing.
rulebasePanel.rbName	The name of the rule base (existing or new).
rulebasePanel.rbPath	The path for the new rule base. There is no default location. You must specify a path.
rulebasePanel.fromRB	If you are creating a new rule base, identify any existing rule bases that you want to import into the new rule base.
bckupERB.backupName	If you want the install to back up an existing rule base before modifying it, specify the name of the backup rule base.
rulebasePanel.backupPath	Specify the directory where the backup rule base should be created.
restartTECQ.restartTEC	Specify whether or not the installer should recycle the event server to implement the changes. Specify Yes to have the installer restart the server; specify No if you want to restart the server yourself after installation is completed.

4. Save the file.

5. Run the following command:

Windows

```
ESync3000Win32.exe -options filename -silent
```

where *filename* is the name of your configuration file.

UNIX

```
ESync3000operating_system.bin -options filename -silent
```

where *operating_system* is the operating system you are installing on (Aix, HP11, Linux, linux390, Solaris). For example, on AIX, run the following command:

```
ESync3000Aix.bin -options filename -silent
```

You must stop and restart the event server for these changes to take effect.

Results

When installation is complete, the results are written to the itm_tec_event_sync_install.log file. On UNIX, this log file is always created in the /tmp directory. For Windows, this file is created in the directory defined by the %TEMP% environment variable. To determine where this directory is defined for the current command-line window, run the following command:

```
echo %TEMP%
```

What to do next

If you specified the monitoring servers in the silent installation configuration file, you might consider deleting that file after installation, for security reasons. The passwords specified in the files are not encrypted.

If you want to define additional monitoring servers (in addition to the one required monitoring server), run the sitconfsvruser.sh command as described in “Defining additional monitoring servers to the event server” on page 822. Repeat this command for each monitoring server.

If you specified your monitoring servers after the installation, you must stop and restart the Situation Update Forwarder process manually. See “Starting and stopping the Situation Update Forwarder process” on page 821 for information.

Perform the following tasks after the installation is finished:

- Install the monitoring agent .baroc files on the event server as described in “Installing monitoring agent .baroc files on the event server” on page 819.
- Configure the monitoring server to forward events to the event server as described in “Configuring your monitoring server to forward events” on page 820.
- If you did not choose to have the rule base updated automatically, update the rule base as described in “Manually importing the event synchronization class files and rule set.”

Manually importing the event synchronization class files and rule set

If you do not want to permit the installation program to modify your rule base, you can choose the manual rule base modification option during the installation.

Then use one of the following methods to manually modify your rule base:

- “Creating a new rule base” on page 817
- “Creating a new rule base and importing an existing rule base into it” on page 817
- “Modifying an existing rule base” on page 818

Before you can run any of the commands in the following sections, you must source your Tivoli environment by running the following command:

Windows Run the following command from a command prompt:

```
C:\Windows\system32\drivers\etc\Tivoli\setup_env.cmd
```

Linux **UNIX** Run the following command from a shell environment:

```
. /etc/Tivoli/setup_env.sh
```

See the *IBM Tivoli Enterprise Console Command and Task Reference* for more information about the **wrb**, **wstopesvr**, and **wstartesvr** commands.

Creating a new rule base

This section outlines the steps for creating a new rule base after you install the event synchronization component.

About this task

Use the following steps to create a new rule base after you install the event synchronization component:

Procedure

1. Create the new rule base by running the following command:

```
wrb -crtrb -path newrb_path newrb_name
```

where *newrb_path* is the path to where you want to create the new rule base, and *newrb_name* is the name for the new rule base.
2. Import the event synchronization class and rule files into the new rule base from the \$BINDIR/TME/TEC/OM_TEC/rules directory created during the installation of the event synchronization component. Run the following commands:

```
wrb -imprbclass path_to_Sentry_baroc_file newrb_name  
wrb -imprbclass path_to_omegamon_baroc_file newrb_name  
wrb -imprbrule path_to_omegamon_rls_file newrb_name  
wrb -imptgtrule omegamon EventServer newrb_name
```
3. Compile and load the new rule base by running the following commands:

```
wrb -comprules newrb_name  
wrb -loadrb newrb_name
```
4. Stop and restart the event server by running the following commands:

```
wstopesvr  
wstartesvr
```

Creating a new rule base and importing an existing rule base into it

This section outlines the steps for creating a new rule base and importing an existing rule base into it.

About this task

Use the following steps to create a new rule base and import an existing rule base into it:

Procedure

1. Create the new rule base by running the following command

```
wrb -crtrb -path newrb_path newrb_name
```

where *newrb_path* is the path to where you want to create the new rule base and *newrb_name* is the name for the new rule base.
2. Import the existing rule base into the new rule base by running the following commands:

```
wrb -cprb -overwrite existing_rbname newrb_name
```

where *existing_rbname* is the name of the existing rule base that you want to import.
3. If the existing rule base is an older rule base, you must upgrade the tec.baroc file to include the TEC_Generic class. Run the following command:

```
perl $BINDIR/TME/TEC/OM_TEC/bin/upg_tec_baroc.pl newrb_name
```
4. If the rule base already contains a Sentry.baroc file, you must upgrade it with the event synchronization event class attributes. Run the following command:

```
perl $BINDIR/TME/TEC/OM_TEC/bin/upg_sentry_baroc.pl
```

5. If the rule base does not contain a Sentry.baroc file, you must import it from the \$BINDIR/TME/TEC/OM_TEC/rules directory created during event synchronization installation. Run the following command:

```
wrb -imprbclass path_to_Sentry_baroc_file newrb_name
```

6. Import the omegamon.baroc and rules file into the rule base from the \$BINDIR/TME/TEC/OM_TEC/rules directory created during event synchronization installation. Run the following commands:

```
wrb -imprbclass path_to_omegamon_baroc_file newrb_name  
wrb -imprbrule path_to_omegamon_rls_file newrb_name  
wrb -imptgtrule omegamon EventServer newrb_name
```

7. Compile and load the new rule base by running the following commands:

```
wrb -comprules newrb_name  
wrb -loadrb newrb_name
```

8. Stop and restart the event server by running the following commands:

```
wstopesvr  
wstartesvr
```

Modifying an existing rule base

This section outlines the steps for modifying an existing rule base to include the class files and rule set for the event synchronization component.

About this task

Use the following steps to modify an existing rule base to include the class files and rule set for the event synchronization component:

Procedure

1. If the existing rule base is an older rule base, you must upgrade the tec.baroc file to include the TEC_Generic class. Run the following command:
2. If the rule base already contains a Sentry.baroc file, you must upgrade it with the event synchronization event class attributes. Run the following command:
3. If the rule base does not contain a Sentry.baroc file, you must import it from the \$BINDIR/TME/TEC/OM_TEC/rules directory created during event synchronization installation. Run the following command:

```
wrb -imprbclass path_to_Sentry_baroc_file newrb_name
```

4. Import the omegamon.baroc and rules file into the rule base from the \$BINDIR/TME/TEC/OM_TEC/rules directory created during event synchronization installation. Run the following commands:

```
wrb -imprbclass path_to_omegamon_baroc_file newrb_name  
wrb -imprbrule path_to_omegamon_rls_file newrb_name  
wrb -imptgtrule omegamon EventServer newrb_name
```

5. Compile and load the new rule base by running the following commands:

```
wrb -comprules newrb_name  
wrb -loadrb newrb_name
```

6. Stop and restart the event server by running the following commands:

```
wstopesvr  
wstartesvr
```

Installing monitoring agent .baroc files on the event server

The monitoring server generates Tivoli Enterprise Console events with classes that are unique to each monitoring agent. Each monitoring agent provides a .baroc file with the Tivoli Enterprise Console classes that are generated by IBM Tivoli Monitoring.

About this task

In order to view this event data in the event console, you must install these monitoring agent .baroc files on the event server.

After you have added application support for each agent to the monitoring server, the monitoring agent .baroc files are located in the following directory:

- **Windows** In the *itm_installdir*\cms\TECLIB directory, where *itm_installdir* is the directory where you installed IBM Tivoli Monitoring.
- **Linux** **UNIX** In the *itm_installdir*/tables/*ms_name*/TECLIB directory, where *itm_installdir* is the directory where you installed IBM Tivoli Monitoring and *ms_name* is the name of the monitoring server.
- **z/OS** users, see *Configuring the Tivoli Enterprise Monitoring Server on z/OS* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm).

Use the following steps to install the monitoring agent .baroc files on the event server:

Procedure

1. Copy the monitoring agent .baroc files from the computer where the monitoring server is installed to a temporary directory on the event server computer (for example, /tmp). The location of the agent .baroc files is described above. Do not copy the om_tec.baroc file; this file contains classes that are duplicates of classes in the omegamon.baroc file.
2. Set up the Tivoli Management Framework environment by running the following command:

Windows Run the following command:

```
C:\WINDOWS\system32\drivers\etc\Tivoli\setup_env.cmd
```

Linux **UNIX** Run the following command from a shell environment:

```
. /etc/Tivoli/setup_env.sh
```

3. For each monitoring agent .baroc file to load into the rule base, run the following command from the same command prompt:

```
wrb -imprbclass /tmp/agent_baroc_file rb_name
```

where:

/tmp/agent_baroc_file

Specifies the location and name of the monitoring agent .baroc file. The example above uses the /tmp directory as the location.

rb_name

Is the name of the rule base that you are using for event synchronization.

4. Compile and load the rule base by running the following commands:

```
wrb -comprules rb_name  
wrb -loadrb rb_name
```

5. Stop and restart the event server by running the following commands:

```
wstopesvr  
wstartesvr
```

Results

When you have loaded each of the agent .baroc files into the rule base and restarted the event server, the event server is ready to receive and correctly parse any events it receives from the monitoring server from one of the installed monitoring agents.

See the *IBM Tivoli Enterprise Console Command and Task Reference* for more information about the **wrb**, **wstopesvr**, and **wstartesvr** commands.

Configuring your monitoring server to forward events

Before the monitoring server forwards any situation events to Tivoli Enterprise Console, you have to enable that forwarding of events.

About this task

Use the following steps to enable event forwarding on your monitoring server.

Note: z/OS users, see *Configuring the Tivoli Enterprise Monitoring Server on z/OS* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm).

For Windows monitoring servers *only*, do the following:

Procedure

1. Open **Manage Tivoli Enterprise Monitoring Services**.
2. Right-click the monitoring server and click **Reconfigure**.
3. On the configuration options window, select Tivoli Event Integration Facility.
4. Click **OK** and **OK**.
5. Complete the following fields on the Event Server: Location and Port Number window and click **OK**:

Server or EIF Probe Location

Type the host name or IP address for the computer where the IBM Tivoli Enterprise Console event server is installed.

Port Number

Type the port number for the event server. If the event server is using port mapping, set this value to 0. If the event server was configured to use a specific port number, specify that number.

To determine the port number that the event server is using, search for the `tec_rcv_agent_port` parameter in the `.tec_config` file in the `$BINDIR/TME/TEC` directory on the event server. If the parameter is commented out with a pound sign (#), the event server is using port mapping. If it is not, the event server is using the port number specified by this parameter.

Note: If you already have policies that contain emitter activities that send events to the Tivoli Enterprise Console, turning on Tivoli Event Integration event forwarding will result in duplicate events. You can deactivate the emitter activities within policies so you do not have to modify all your policies when you activate Tivoli Event Integration Facility forwarding by specifying **Disable Workflow Policy/Tivoli Emitter Agent Event Forwarding** when you enable forwarding using the Event Integration Facility.

Using policies gives you more control over which events are sent and you may not want to lose this granularity. Moreover, it is likely the policies that are invoking the TEC emitter are doing little else. If you deactivate these activities, there is no point in running the policy. You may delete policies that are

longer required, instead of disabling them. Note that events forwarded via the TEC event emitter are not eligible for event synchronization (that is, changes to these events on the TEC side will not be sent back to the monitoring server).

What to do next

For Linux and UNIX monitoring servers: You configured the TEC Server and TEC Port information for the Linux/UNIX monitoring server during installation, if you installed the monitoring server using the configuration instructions in this installation guide. However, if you did not configure this information, see for the procedure.

Controlling event forwarding

The `om_tec.config` file controls the forwarding of events to IBM Tivoli Enterprise Console.

About this task

The `om_tec.config` file contains the following parameter:

BufferFlushRate=events_per_minute

Specifies the number of events that are sent per minute when the adapter has reestablished its connection to Tivoli Enterprise Console. After the adapter has recovered the lost connection, if there are events in the buffer, the events are sent at this rate per minute. The default value is 0, meaning all events are sent in one burst.

If your environment can have a large number of open situation events, you might want to adjust this parameter to control the rate at which events are sent to the event server. To edit this file and change this parameter:

- **Windows** :
 1. Open Manage Tivoli Enterprise Monitoring Services.
 2. Right-click Tivoli Enterprise Monitoring Server, and click Advanced → Edit EIF Configuration.
 3. Once you have completed your reconfiguration, recycle the Tivoli Enterprise Portal Server and each Tivoli Enterprise Portal client.
- **Linux** **UNIX** :
 1. Edit file `install_dir/tables/hostname/TECLIB/ om_tec.config`, where `install_dir` is your Tivoli Monitoring installation directory and `hostname` is the name of the host running this monitoring server.
 2. Once you have saved your configuration updates, recycle the Tivoli Enterprise Portal Server and each Tivoli Enterprise Portal client.

Starting and stopping the Situation Update Forwarder process

To send event updates to a monitoring server, you must start the Situation Update Forwarder. This process is started automatically when the event server starts.

To start the process manually, change to the `$BINDIR/TME/TEC/OM_TEC/bin` directory (where `$BINDIR` is the location of the Tivoli Management Framework installation) and run the following command:

Windows :
`startSUF.cmd`

UNIX :
`startSUF.sh`

To stop the process, run the following command:

Windows :
stopSUF.cmd

UNIX :
stopSUF.sh

On Windows, you can also start and stop the Tivoli Situation Update Forwarder service to start or stop the forwarding of event updates. You can start and stop this service either from the Windows Service Manager utility or with the following commands:

```
net start situpdate
net stop situpdate
```

Changing the configuration of the event synchronization component on the event server

If you want to change any of the settings for the event synchronization on the event server, use the **sitconfig.sh** command.

You have two options for running this command:

- Manually modify the configuration file for event synchronization (named situpdate.conf by default and located in the /etc/TME/TEC/OM_TEC directory on UNIX and Linux, and the %SystemDrive%\Program Files\TME\TEC\OM_TEC\etc directory on Windows), and then run the following command:
`sitconfig.sh update filename=config_filename`
- Run the **sitconfig.sh** command directly, specifying only those settings that you want to change. See the *IBM Tivoli Monitoring Command Reference* for the full syntax of this command.

After you change the configuration of the event synchronization, you must manually stop and restart the Situation Update Forwarder process. See “Starting and stopping the Situation Update Forwarder process” on page 821 for information.

Defining additional monitoring servers to the event server

This section outlines the steps for defining additional monitoring servers to the event server.

About this task

To add additional monitoring servers to the list that can receive event status updates from Tivoli Enterprise Console, run the sitconfsvruser.sh command as follows:

Procedure

1. Source your Tivoli environment by running the following command:
 - **Windows** Run the following command from a command prompt:
`C:\Windows\system32\drivers\etc\Tivoli\setup_env.cmd`
 - **Linux** **UNIX** Run the following command from a shell environment:
`. /etc/Tivoli/setup_env.sh`
2. **Windows** Type BASH to invoke the bash shell.
3. Change to the \$BINDIR/TME/TEC/OM_TEC/bin directory (where \$BINDIR is the location of the Tivoli Management Framework installation) and enter the following command:
`sitconfsvruser.sh add serverid=server userid=user password=password`

where:

server Is the fully qualified host name of the monitoring server.

user Is the user ID to access the computer where the monitoring server is running.

password
Is the password to access the computer.

What to do next

Repeat this command for each monitoring server.

You can also delete monitoring servers. See the *IBM Tivoli Monitoring Command Reference* for the full syntax of this command.

After you change the configuration of the event synchronization, you must manually stop and restart the Situation Update Forwarder process. See “Starting and stopping the Situation Update Forwarder process” on page 821 for information.

Customizing event status processing behavior when agent switching is used or the agent goes offline

You can add variables to the monitoring server's environment file to customize the behavior of event status processing.

The variables in Table 152 on page 824 can be added to the monitoring server's environment file to customize the behavior of event status processing when agent switching is used or when the agent goes offline. The first two variables help ensure that events are not closed by the agent's primary monitoring server after the agent has switched to its secondary monitoring server. You can find the monitoring server's environment file in these locations:

- **Windows**

ITM_HOME\cms\KBBENV
For example: C:\IBM\ITM\cms\KBBENV

- **Linux** **UNIX**

ITM_HOME/config/tems_hostname_ms_tems_name.config
For example: /opt/IBM/ITM/config/edinburg_ms_labtems.config

For Linux/UNIX systems, you must add the variables to the .config and the .ini files. The name and location of the .ini file is ITM_HOME/config/ms.ini.

- **z/OS**

&shilev.&rtename.RKANPARU(KDSENV)
For example: ITM.SYP1.RKANPARU(KDSENV)

Note: The &shilev and &rtename are variables that correspond to high level qualifiers of the RKANPARU(KDSENV) partitioned dataset. These variables can take 1 to 8 characters.

You must recycle the Tivoli Enterprise Monitoring Server after modifying the environment file for your changes to be picked up.

Table 152. Variables to customize the behavior of event status processing when agent switching is used.

Variable	Architecture Type	Details	Administrator
IRA_MIN_NO_DATA_WAIT_TIME	Unidirectional and bidirectional	<p>The minimum time to wait before the monitoring server closes a situation event. This parameter is defined in number of seconds. The default value is zero.</p> <p>By default, after an agent is disconnected from a Tivoli Enterprise Monitoring Server, situations already open will remain open for three situation polling intervals. For example, take two sampled situations, S1 and S2, with intervals of 30 seconds and 15 minutes respectively. Both situations are open when the agent loses connection. Situation S1 closes after at least one minute and 30 seconds. Situation S2 closes after at least 45 minutes. With agent switching, if a situation closes too soon it might generate duplicate events because the agent did not have sufficient time to connect to the backup monitoring server before the primary server closes the original event. This is particularly true for situations with very short polling intervals.</p> <p>In such a scenario you can use the IRA_MIN_NO_DATA_WAIT_TIME variable to set the minimum wait time before a situation is closed. Using the example above, if IRA_MIN_NO_DATA_WAIT_TIME is set to 300 (5 minutes), S1 will close after 5 minutes not 90 seconds. S2 is unaffected and will close after 45 minutes as before.</p> <p>Note: You should set this variable in the environment file for all of your monitoring servers.</p>	IBM Tivoli Monitoring
CMS_SIT_TIME_VALIDATION	Unidirectional and bidirectional	<p>Valid entries are Y or N. The default is N.</p> <p>By default, the monitoring server handles situation events on a first-come-first-serve basis. In a scenario where agent switching is enabled an agent might send events through two different monitoring servers. The events that arrive first might not necessarily be the earlier events if one of the monitoring servers encountered connection issues. This generally has little impact on situation event processing, except when a monitoring server is shutdown and some situations might be closed prematurely even though the agent is already connected to a different monitoring server.</p> <p>You must perform two actions to avoid this scenario:</p> <ol style="list-style-type: none"> 1. All monitoring server hosts should synchronize time, preferable through Internet Time Protocol (ITP) clients. 2. You should add the CMS_SIT_TIME_VALIDATION=Y variable to all monitoring server environment files. This switches the LCLTMSTMP column in situation events to use UTC time instead of local time, which is then used to determine event order. 	IBM Tivoli Monitoring

Table 152. Variables to customize the behavior of event status processing when agent switching is used. (continued)

Variable	Architecture Type	Details	Administrator
CMS_SIT_CHECK_NODESTS	Unidirectional and bidirectional	<p>Valid entries are Y or N. The default is N.</p> <p>This variable is only applicable for users of event integration with Tivoli Enterprise Console and Omnibus. When the CMS_SIT_CHECK_NODESTS variable is set to Y in the environment file, the hub monitoring servers check the agent status whenever a close status update event is forwarded to Netcool/Omnibus. The CMS_SIT_CHECK_NODESTS variable should only be added to the hub monitoring server environment file. If the agent is offline the close status update event is tagged with a special OFFLINE indicator in the situation_eventdata EIF slot.</p> <p>If you do not want events to be closed in Netcool/OMNIBus when an agent goes offline, you can customize the EIF probe rules to ignore close events where the situation_status slot is set to N and the situation_eventdata EIF slot is set to OFFLINE. See “Customizing the rules file” on page 762 for details on how to add customizations to the EIF probe rules. You should also consider setting the IRA_MIN_NO_DATA_WAIT_TIME environment variables described in this table so that close status events are not sent to Netcool/OMNIBus until after the agent offline condition has been detected. (If the default agent heartbeat interval is used, it can take between 10 to 20 minutes before the monitoring server detects that the agent is no longer online.)</p> <p>Note: You should only set the CMS_SIT_CHECK_NODESTS variable in the environment file of your hub monitoring server.</p>	IBM Tivoli Monitoring

Changing the TCP/IP timeout setting on your event server

If the Situation Update Forwarder cannot reach a monitoring server to send an update, depending on the TCP/IP settings for the computer where your event server is running, you might have to wait up to 15 minutes before the Situation Update Forwarder tries to connect to the monitoring server again. This might occur if your event server is running on an AIX, Solaris, or HP-UX computer.

Use the following steps to change the TCP/IP timeout value for your computer.

On AIX, run the following command:

```
no -o tcp_keepinit=timeout_value
```

where *timeout_value* is the length of the timeout period, in half seconds. To configure a timeout of 30 seconds, set the *timeout_value* value to 60.

On Solaris and HP-UX, run the following command:

```
ndd -set /dev/tcp tcp_ip_abort_cinterval timeout_value
```

where *timeout_value* is the length of the timeout period, in milliseconds. To configure a timeout of 30 seconds, set the *timeout_value* value to 30000.

Upgrading to Tivoli Event Synchronization version 3.0.0.0

Upgrading replaces the *omegamon.rls* file. Any changes you have made to this file will be lost. However, a backup file named *omegamon.rls.bac* is created so you can recover your changes.

You must load the noncurrent rule base and restart the Tivoli Enterprise Console event server after updating either the noncurrent or the current rule base.

Note: If you have already installed IBM Tivoli Monitoring and Tivoli Event Synchronization, you must upgrade to version 3.0.0.0.

You install the upgrade from the IBM Tivoli Monitoring installation media.

There are three methods for upgrading event synchronization:

- “Upgrading from a wizard”
- “Upgrading from the command-line” on page 827
- “Upgrading from the command-line using a silent installation” on page 829

If you have multiple rule bases that are using IBM Tivoli Monitoring and Tivoli Event Synchronization, you can run the upgrade installation to update each rule base. After you finish the first rule base, restart the upgrade installer and supply the targeted next rule base you want to update.

Note: You cannot uninstall just the upgrade. You must uninstall the entire product. For instructions on uninstalling, see “Uninstalling the event synchronization component” on page 940.

Upgrading from a wizard

This section outlines the steps for upgrading event synchronization from a wizard.

About this task

Use the following steps to upgrade event synchronization from the installation wizard:

Procedure

1. On the host of the event server, launch the event synchronization upgrade installation:

Windows Double-click the *ESUpgrade30Win32.exe* file in the *tec* subdirectory on the *IBM Tivoli Monitoring Tools* DVD or DVD image.

Linux **UNIX** Change to the *tec* subdirectory on the *IBM Tivoli Monitoring Tools* DVD or DVD image and run the following command:

```
ESUpgrade30operating_system.bin
```

where *operating_system* is the operating system you are installing on (*aix*, *HP11*, *Linux*, *linux390*, or *Solaris*). For example, run the following command on an AIX computer:

```
ESUpgrade30Aix.bin
```

2. Click **Next** on the Welcome window.
3. Select **I accept the terms in the license agreement** and click **Next**. You are presented with the option of having the installer automatically update the rule base for you.
4. Select whether you want to have the installer update the rule base or whether you update it manually after the installation is complete, then click **Next**. If you chose to manually update the rule base, you will see a progress indicator for the installation. Proceed to step 9. If you chose to have the installer automatically update the rule base, the data collection window shown in Figure 159 on page 827

829 is displayed:
Proceed to step 5.

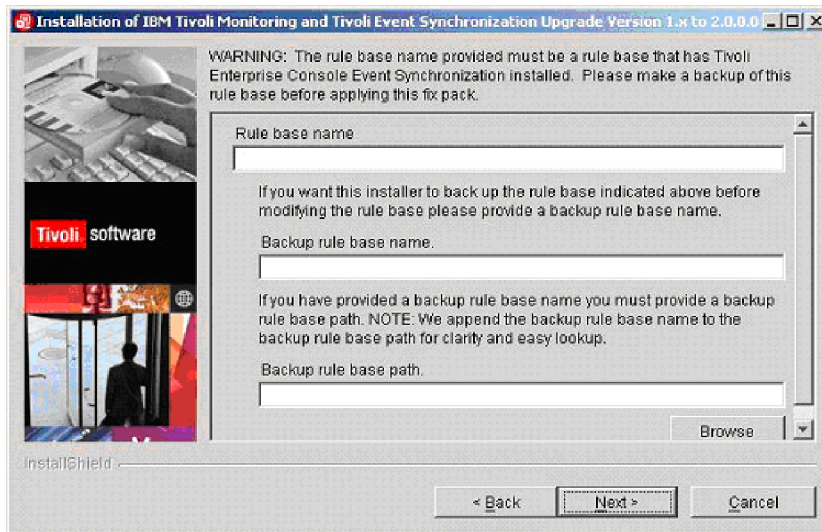


Figure 159. Upgrade data collection window

5. Specify the name of the rule base to be upgraded. The rule base must be one that has event synchronization previously installed.
6. If you want the installer to back up the rule base before it is modified, specify a name and a path for the backup rule base.
7. Click **Next** to continue. A window is displayed that summarizes the information you entered.
8. If the information is correct, click **Next** to proceed with the installation. If the information is not correct, click **Back** and correct the fields as necessary; then click **Next** and **Next** again to proceed. A progress indicator shows the progress of the installation and configuration.
9. When the installation completes successfully, you will see a message that reminds you to restart the TEC server. If the updated rule base is not the currently loaded rule base, you are reminded to load the rule base and restart the server. Click **OK** to dismiss the message. A window is displayed that reminds you to restart the Tivoli Enterprise Console server.
10. If you want the installer to restart the server for you, check Restart the Tivoli Enterprise Console server to make changes effective; then click **Next**. If you do not want the installer to restart the server, leave the option unchecked and click **Next**.
11. Click **Finish** to exit the installer.

What to do next

Important:: If you chose the manual update option, you must copy the files in \$BINDIR/TME/TEC/OM_TEC/rules directory to the rule base, recompile and reload the rule base, and restart the Tivoli Enterprise Console. See “Manually importing the event synchronization class files and rule set” on page 816 for the commands to use to do this.

Upgrading from the command-line

This section outlines the steps for upgrading event synchronization from the command-line on your event server.

About this task

Use the following steps to upgrade event synchronization from the command-line on your event server:

Procedure

1. Change to the `tec` directory on the IBM Tivoli Monitoring Tools DVD or DVD image.
2. Run the following command to launch the installation:

Windows

```
ESUpgrade30Win32.exe -console
```

Linux

UNIX

```
ESUpgrade30operating_system.bin -console
```

where *operating_system* is the operating system you are installing on (Aix, HP11, Linux, linux390, or Solaris). For example, run the following command on an AIX computer:

```
ESUpgrade30Aix.bin -console
```

The following prompt is displayed:

Press 1 for Next, 3 to Cancel or 4 to Redisplay [1]

3. Type 1 to start the installation and press Enter.

The following prompt is displayed:

Software Licensing Agreement:

Press Enter to display the license agreement on your screen. Please read the agreement carefully before installing the Program. After reading the agreement, you will be given the opportunity to accept it or decline it. If you choose to decline the agreement, installation will not be completed and you will not be able to use the Program.

4. Press Enter to display the software license agreement.
5. Type 1 and press Enter to accept the license.

The following prompt is displayed:

```
[X] 1 - Automatically install rules and classes (recommended)
[ ] 2 - Manually install rules and classes (advanced users)
```

To select an item enter its number, or 0 when you are finished: [0]

6. If you want to have the installer automatically install the rules and classes, enter 1. The following prompt is displayed:

Rule base Name []

Continue with step 7. If you manually install the rules and classes, enter 2 and proceed to step 11.

7. Type 1 and press Enter to continue.
8. Type the name of the rule base to upgrade then press Enter.

The rule base must be one in which event synchronization was previously installed.

The following prompt is displayed:

Backup rule base name. []

9. If you want the installer to back up the rule base before modifying it, type a name for the backup rule base, then press Enter. If you do not want the installer to create a backup rule base, press Enter without typing any name. If you provide a name for the backup rule base, you are prompted for the path for the backup rule base:

If you have provided a backup rule base name you must provide a backup rule base path. NOTE: We append the backup rule base name to the backup rule base path for clarity and easy look-up.

Backup rule base path. []

10. Type the path and press Enter. The following prompt is displayed:
Press 1 for Next, 2 for Previous, 3 to Cancel, or 4 to Redisplay []
11. Type 1 and press Enter to continue.

The following prompt is displayed:

IBM Tivoli Monitoring

Press 1 for Next, 2 for Previous, 3 to Cancel, or 4 to Redisplay [1]

12. Type 1 and press Enter to continue.

The event synchronization is installed. The following prompt is displayed:

Installation and Configuration has completed. Please stop and restart the Tivoli Enterprise Console Server.

Press 1 for Next, 2 for Previous, 3 to Cancel, or 4 to Redisplay [1]

13. Type 1 and press Enter to continue.

The following prompt is displayed:

Installation and configuration has completed.
Please restart the Tivoli Enterprise Console server for the changes to take effect.

Mark appropriately below to restart the Tivoli Enterprise Console server.
[] 1 - Restart the Tivoli Enterprise Console server to make changes effective

To select an item enter its number, or 0 when you are finished: [0]

If you did not choose to have the installer automatically update the rule base, you will not be offered the option to restart the Tivoli Enterprise Console automatically.

14. If you want the installer to stop and restart the Tivoli Enterprise Console server, type 1 and press Enter. If you want to stop and restart the console yourself, type 0 and press Enter. The following prompt is displayed:

Press 3 to Finish, or 4 to Redisplay [1]

15. Type 3 to finish and press Enter. You must stop and restart the event server for these changes to take effect.

What to do next

Important: If you chose the manual update option, you must copy the files in \$BINDIR/TME/TEC/OM_TEC/rules directory to the rule base, recompile and reload the rule base, and restart the Tivoli Enterprise Console. See “Manually importing the event synchronization class files and rule set” on page 816 for the commands to use to do this.

Upgrading from the command-line using a silent installation

This section outlines the steps for upgrading event synchronization from the command-line using a silent installation.

About this task

If you have already installed Fix Pack 1, be aware that if you define the rulebasePanel.rbName, this silent installation will update the rule base. If you do not want to apply the updated rules file to a rule base leave the rulebasePanel.rbName commented out (do not delete the pound signs ### preceding the parameter).

Use the following steps to install the event synchronization using a silent installation from the command-line on your event server. This installation method runs silently, so you will not see status messages during the actual installation.

Procedure

1. Change to the `tec` directory on the IBM Tivoli Monitoring Tools DVD or DVD image.
2. Run the following command to generate the configuration file:

Windows

```
ESUpgrade30Win32.exe -options-template filename
```

where *filename* is the name of the configuration file to create, for example, `es_silentinstall.conf`.

UNIX

```
ESUpgrade30operating_system.bin -options-template filename
```

where *operating_system* is the operating system you are installing on (Aix, HP11, Linux, linux390, or Solaris). For example, run the following command on an AIX computer:

```
ESUpgrade30Aix.bin -options-template filename
```

3. Edit the output file to specify the following values.

rulebasePanel.rbInstallType

Specifies whether the installer will automatically update a specified rule base, or if a rule base must be manually modified after installation is complete. Specify either automatic or manual. If manual is specified, all other rulebasePanel.* options are ignored.

rulebasePanel.rbName

The name of the rule base to be updated.

rulebasePanel.backupName

The name of the backup rule base.

rulebasePanel.backupPath

The path for the backup rule base.

restartTEC.startTEC

Indicates that you want the installer to restart the Tivoli Enterprise Console event server.

Note:

- a. Remove the pound signs (###) from the beginning of any value that you want to specify.
 - b. Do not enclose any values in quotation marks (").
 - c. Specify the following value only if you want the indicated rule base updated:
`rulebasePanel.rbName`
 - d. If you specify values, ensure that the value you specify meets the minimum required values. Otherwise, the installation stops and an error is written to the log file.
4. Save the file.
 5. Run the following command:

Windows

```
ESUpgrade30Win32.exe -options filename -silent
```

where *filename* is the name of your configuration file.

UNIX

```
ESUpgrade30operating_system.bin -options filename -silent
```

where *operating_system* is the operating system you are installing on (Aix, HP11, Linux, linux390, Solaris). For example, on AIX, run the following command:

```
ESUpgrade30Aix.bin -options filename -silent
```

Results

The rule base that is updated during silent installation is made current.

What to do next

Important: If you chose the manual update option, you must copy the files in \$BINDIR/TME/TEC/OM_TEC/rules directory to the rule base, recompile and reload the rule base, and restart the Tivoli Enterprise Console. See “Manually importing the event synchronization class files and rule set” on page 816 for the commands to use to do this.

Appendix A. Installation worksheets

Use worksheets to gather the information you need during the installation of the IBM Tivoli Monitoring components.

- “Windows hub monitoring server worksheet”
- “Windows Tivoli Enterprise Monitoring Automation Server worksheet” on page 834
- “Linux or UNIX hub monitoring server installation worksheet” on page 834
- “Linux or UNIX Tivoli Enterprise Monitoring Automation Server worksheet” on page 835
- “Windows remote monitoring server worksheet” on page 835
- “Linux or UNIX remote monitoring server installation worksheet” on page 836
- “Windows portal server worksheet” on page 836
- “Linux portal server worksheet” on page 837
- “Generic Windows monitoring agent worksheet” on page 838
- “Generic Linux or UNIX monitoring agent worksheet” on page 838
- “Windows portal desktop client worksheet” on page 838
- “Linux portal desktop client worksheet” on page 839
- “Monitoring server communications protocol details worksheet” on page 839

Windows hub monitoring server worksheet

Gather information for your installation of the hub monitoring server on a Windows computer.

Host name of computer	
IP address	
IBM Tivoli Monitoring installation directory	
Encryption key	Must use the same encryption key on all IBM Tivoli Monitoring components. Do not specify the following characters: & ampersand pipe ' single quote = equal sign \$ dollar sign In addition, do not specify double-byte (DBCS) characters.
Agents to install on this computer	
Agents to add to the deployment depot (plus non-agent bundles, if your site uses them)	
Program folder name	
Monitoring server name	
Communications protocol details	See “Monitoring server communications protocol details worksheet” on page 839.
Agents for which to add application support data	

Windows Tivoli Enterprise Monitoring Automation Server worksheet

Gather information for your installation of the Tivoli Enterprise Monitoring Automation Server on a Windows computer.

A locally installed hub monitoring server is a prerequisite for the Tivoli Enterprise Monitoring Automation Server.	See “Windows hub monitoring server worksheet” on page 833.
Installation directory	
Encryption key for the monitoring server	Must specify the same encryption key on all IBM Tivoli Monitoring components.
Registry Services provider connection information	The HTTP protocol, hostname, and port number for clients to connect to the automation server.
Registry Services provider registry User ID	
Registry Services provider registry Password	
HTTP protocol that clients should use when connecting to the Tivoli Enterprise Monitoring Automation Server	
Port number that clients should use when connecting to the Tivoli Enterprise Monitoring Automation Server	

Linux or UNIX hub monitoring server installation worksheet

Gather information for the installation of the hub monitoring server on a Linux or UNIX computer.

Host name of computer	
IP address	
IBM Tivoli Monitoring installation directory	
Encryption key	Must use the same encryption key on all IBM Tivoli Monitoring components. Do not specify the following characters: & ampersand pipe ' single quote = equal sign \$ dollar sign In addition, do not specify double-byte (DBCS) characters.
Monitoring server name	
Communications protocol details	See “Monitoring server communications protocol details worksheet” on page 839.
KDC_PARTITION	
NIC interface name (Optional Primary Network Name)	
Agents to install on this computer	
Agents for which to add application support data	
Agents to add to the deployment depot (plus non-agent bundles, if your site uses them)	

Linux or UNIX Tivoli Enterprise Monitoring Automation Server worksheet

Gather information for your installation of the Tivoli Enterprise Monitoring Automation Server on a Linux or UNIX computer.

A locally installed hub monitoring server is a prerequisite for the Tivoli Enterprise Monitoring Automation Server.	See “Linux or UNIX hub monitoring server installation worksheet” on page 834.
Installation directory	
Encryption key for the monitoring server	Must specify the same encryption key on all IBM Tivoli Monitoring components.
Registry Services provider connection information	The HTTP protocol, hostname, and port number for clients to connect to the automation server.
Registry Services provider registry User ID	
Registry Services provider registry Password	
HTTP protocol that clients should use when connecting to the Tivoli Enterprise Monitoring Automation Server	
Port number that clients should use when connecting to the Tivoli Enterprise Monitoring Automation Server	

Windows remote monitoring server worksheet

Gather information for the installation of the remote monitoring server on a Windows computer.

Host name of computer	
IP address	
IBM Tivoli Monitoring installation directory	
Encryption key	Must use the same encryption key on all IBM Tivoli Monitoring components. Do not specify the following characters: & ampersand pipe ' single quote = equal sign \$ dollar sign In addition, do not specify double-byte (DBCS) characters.
Agents to install on this computer	
Agents to add to the deployment depot (plus non-agent bundles, if your site uses them)	
Program folder name	
Monitoring server name	
Agents for which to add application support data	
Hub monitoring server name	
Hub monitoring server host name	
Hub monitoring server communications protocol details	See “Monitoring server communications protocol details worksheet” on page 839.

Linux or UNIX remote monitoring server installation worksheet

Gather information for the installation of the remote monitoring server on a Linux or UNIX computer.

Host name of computer	
IP address	
IBM Tivoli Monitoring installation directory	
Encryption key	Must use the same encryption key on all IBM Tivoli Monitoring components. Do not specify the following characters: & ampersand pipe ' single quote = equal sign \$ dollar sign In addition, do not specify double-byte (DBCS) characters.
Monitoring server name	
KDC_PARTITION	
NIC interface name (Optional Primary Network Name)	
Agents for which to add application support data	
Hub monitoring server host name	
Hub monitoring server communications protocol details	See "Monitoring server communications protocol details worksheet" on page 839.
Agents to add to the deployment depot (plus non-agent bundles, if your site uses them)	

Windows portal server worksheet

Gather information for the installation of the portal server on a Windows computer.

Installation location	
Encryption key used on the hub monitoring server	Must specify the same encryption key on all IBM Tivoli Monitoring components.
Program folder	
Host name of the computer where you are installing the portal server	
Portal server database administrator ID	
Portal server database administrator password	
Portal server database user ID (default = TEPS)	
Portal server database user password	
Warehouse database administrator ID	
Warehouse database administrator password	
Warehouse database user ID (default = ITMUser)	
Warehouse database user password	
Warehouse data source name (default = ITM Warehouse)	
Warehouse database name	
Hub monitoring server host name	

Hub monitoring server communications protocol details	See “Monitoring server communications protocol details worksheet” on page 839.
LDAP type	
LDAP base	
LDAP DN base entry	
LDAP bind ID	
LDAP password	
LDAP port number	
LDAP hostname	
Authorization Policy Server protocol	
Authorization Policy Server hostname or IP address	
Authorization Policy Server port number	
Authorization Policy Server user ID	
Authorization Policy Server password	

Linux portal server worksheet

Gather information for the installation of the portal server on a Linux computer.

Installation location	
Encryption key for the hub monitoring server	Must specify the same encryption key on all IBM Tivoli Monitoring components.
Host name for the hub monitoring server	
Hub monitoring server communications protocol details	See “Monitoring server communications protocol details worksheet” on page 839.
NIC interface name (Primary Optional Network Name)	
DB2 for Linux, UNIX and Windows instance name (default = db2inst1)	
DB2 for Linux, UNIX and Windows administrator ID (default = db2inst1)	
DB2 for Linux, UNIX and Windows administrator password	
Portal server database name (default = TEPS)	
Portal server database user (default = itmuser)	
Portal server database user password	
Warehouse database name (default = WAREHOUS)	
Warehouse database user (default = itmuser)	
Warehouse database user password	
LDAP type	
LDAP base	
LDAP DN base entry	
LDAP bind ID	
LDAP password	
LDAP port number	

LDAP hostname	
Authorization Policy Server protocol	
Authorization Policy Server hostname or IP address	
Authorization Policy Server port number	
Authorization Policy Server user ID	
Authorization Policy Server password	

Generic Windows monitoring agent worksheet

Gather information for the installation of a monitoring agent on a Windows computer. Depending on the agent you are installing, you might need additional information to configure the agent. See the agent user's guide for more information.

Installation directory	
Encryption key for the hub monitoring server	Must specify the same encryption key on all IBM Tivoli Monitoring components.
Agents to install	
Program folder name	
Monitoring server host name	
Monitoring server communications protocol details	See "Monitoring server communications protocol details worksheet" on page 839.

Generic Linux or UNIX monitoring agent worksheet

Gather information for the installation of a monitoring agent on a Linux or UNIX computer. Depending on the agent you are installing, you might need additional information to configure the agent. See the agent user's guide for more information.

Installation directory	
Encryption key for the hub monitoring server	Must specify the same encryption key on all IBM Tivoli Monitoring components.
Agents to install	
Agent product code or codes	
Monitoring server host name	
Monitoring server communications protocol details	See "Monitoring server communications protocol details worksheet" on page 839.
KDC_PARTITION	
NIC interface name (Optional Primary Network Name)	
root user password	
User group name	
Optional user name	

Windows portal desktop client worksheet

Gather information for the installation of the portal desktop client on a Windows computer.

Installation directory	
------------------------	--

Encryption key for the monitoring server	Must specify the same encryption key on all IBM Tivoli Monitoring components.
Program folder name	
Portal server host name	
Monitoring server host name	
Monitoring server communications protocol details	See "Monitoring server communications protocol details worksheet."

Linux portal desktop client worksheet

Gather information for the installation of the portal desktop client on a Windows computer.

Installation directory	
Encryption key for the monitoring server	Must specify the same encryption key on all IBM Tivoli Monitoring components.
Instance name for the portal server	
Portal server host name	

Monitoring server communications protocol details worksheet

Gather the communications protocol details for your hub and remote monitoring servers.

IP.UDP Settings
Host name or IP address
Port number or port pools
IP.PIPE Settings
Host name or IP address
Port number
IP.SPIPE Settings
Host name or IP address
Port number
SNA Settings
Network name
LU name
LU 6.2 LOGMODE
TP name
Local LU alias

Appendix B. Performing a silent installation of IBM Tivoli Monitoring

This appendix provides information about installing IBM Tivoli Monitoring using the silent installation method. This method of installation is useful for advanced users who supply installation information once through a response file, instead of repeatedly through an installation wizard.

You might run through the installation wizard one time to determine the values that you need to set for your monitoring needs and then use silent installation to install the rest of your environment. For more information about installing through the installation wizard, see Chapter 10, “Installing IBM Tivoli Monitoring,” on page 193.

The following table outlines the steps for performing a silent installation of IBM Tivoli Monitoring.

Table 153. Installation and configuration steps

Step	Where to find detailed information
Assess your monitoring needs to determine the best deployment of IBM Tivoli Monitoring components.	Chapter 7, “Preparing for installation,” on page 125
Ensure you have the required hardware and software.	Chapter 3, “Hardware and software requirements,” on page 75
Gather any information required for successful installation (such as DB2 user information and security specifications).	“Specific information to have ready” on page 126
Run the silent installation.	“Creating and using a Windows response file” “Performing a silent installation on a Linux or UNIX computer” on page 848
Install application support files on the monitoring server, portal server, and portal desktop client.	“Installing and enabling application support” on page 275
Start the portal client to verify that you can view the monitoring data.	“Starting the Tivoli Enterprise Portal client” on page 366

For information on performing a silent uninstallation, see “Uninstalling components and agents silently” on page 938.

Creating and using a Windows response file

A sample Windows silent installation response file is provided on the product installation media.

About this task

Use the following steps to edit that response file as appropriate for your environment:

Procedure

1. Locate on the product installation media the appropriate silent_*.txt file, as dictated by the Tivoli Management Services component for which you're building a response file.

silent_server.txt

for a server image

silent_agent.txt

for an agent image

silent_WIA64.txt

for an agent image for 64-bit Windows Itanium

For example, when installing an agent silently you must use the appropriate response file located in ITM_build\Windows\Deploy.

2. Copy this file to a temporary directory on your system.
3. Open your copy of the silent_*.txt file in a text editor.
4. Change the parameters as appropriate for your environment. The silent_*.txt file contains descriptions of all the parameters, including directions on how to use them.

Complete all of the steps listed in the file. Each line of the file must be either a comment (containing a semi-colon in column one) or a meaningful statement that starts in column one.

Note: If you want to use the TCP/IP protocol, make sure to specify "IP.UDP." If you specify "TCP/IP," IP.PIPE is used by default.

Attention: Do not modify any other files supplied with the installation (for example, the SETUP.ISS file).

5. Save the file and close the editor.
6. Run the silent installation using one of the following methods:
 - "Running the silent installation from the command-line with parameters" on page 844
 - "Running the silent installation using SMS" on page 845

Results

Note:

1. A silentInstall.cmd script has been added to the Agents DVD. To install this agent you need to run this script:

```
silentInstall.cmd
```

To install the agent in a different directory than the default one (CANDLE_HOME), use the -h option:

```
silentInstall.cmd -h directory
```

If this directory name contains spaces, make sure you enclose the name in quotation marks:

```
silentInstall.cmd -h "directory_with_spaces"
```

To review the usage of the silentInstall.cmd file, enter `silentInstall.cmd -?`.

2. To silently install the tacmd command interface (component KUE) so you can continue running commands based on the previous CLI, uncomment this line in the FEATURES sections of your response file:

```
KUEWICMA= Tivoli Enterprise Services User Interface
```
3. If the installation fails for any reason, a log file, "Abort IBM Tivoli Monitoring *date time*.log," is created to document the problem. If the installation fails before reading in the installation location, the log file is written to the Windows boot drive, typically the C:\ drive. If the installation fails after reading the installation location, the log file is written to an \install subdirectory in the installation directory. For example, if you use the default installation directory, the log file is written to the C:\ibm\itm\install\ITM directory.

Automatically creating agent response files on Windows

You can have Tivoli Monitoring create the response files for you after configuring an agent.

About this task

This new feature vastly simplifies and speeds up the creation and customization of response files for agent installation and deployment. It also reduces the likelihood of user error when specifying their contents. The resulting response files can be used to either install or deploy similar agents across your environment.

Note: The automatic generation of response files does not apply to multi-instance agents or to server components.

The agent must be successfully installed and configured prior to generating the response file. To have Tivoli Monitoring generate a response file, follow these steps:

Procedure

1. From the Manage Tivoli Enterprise Monitoring Services screen, right-click the agent whose configuration information you want saved, and select the **Advanced** → **Utilities** → **Generate Response Files** option from the pop-up menu. See Figure 160.

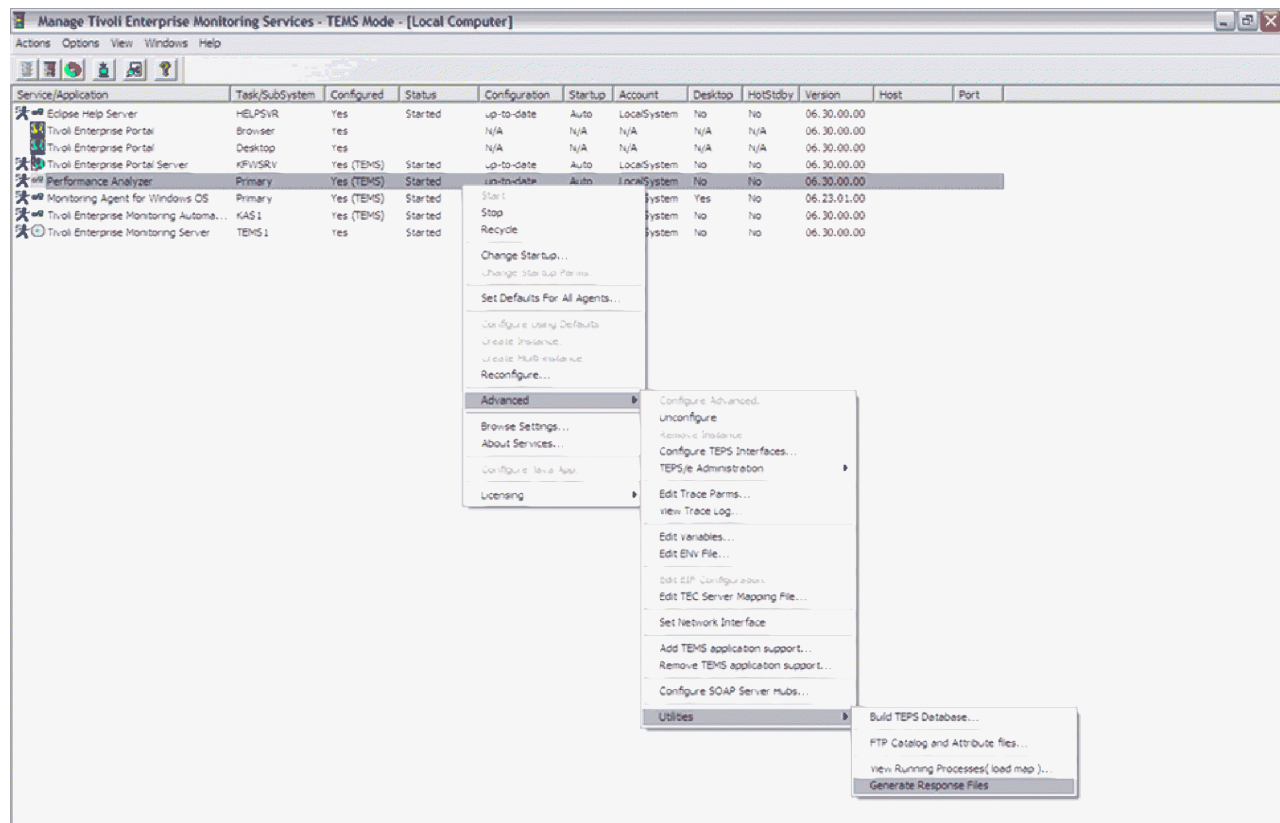


Figure 160. The Generate Response Files option

Note: This option does not appear if the agent has not yet been configured.

2. A window opens where you can specify the path into which you want the generated response files written; the default directory is `ITM_Install_Home\response`.

The newly generated response files are named `silent_install_pc.txt` and `silent_deploy_pc.txt`, where `pc` is the product code for the agent whose configuration parameters you want saved. See Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883 for the list of product codes. If the directory you specify already contains files with these names, you are prompted to either replace them or specify new names.

Results

The installation and remote-deployment response files are created.

What to do next

When the installation and remote-deployment response files are available, you can use them to either install or remotely deploy identical agents across your IBM Tivoli Monitoring network:

1. Edit the response file, and supply missing security parameters such as passwords and encryption keys.

Note: For security reasons, the IBM Tivoli Monitoring encryption key entered during installation is not saved in the generated response file, and password parameters are blanked out. You must supply these values yourself.

2. Install the agent.
 - If you install it locally on the target computer, perform a silent installation using the generated response file for installation: `silent_install_pc.txt`.
 - If instead you remotely deploy the new agent, use the generated response file for remote deployment: `silent_deploy_pc.txt`.

When the silent installation completes successfully, the new agent is installed and configured with identical settings as the first one.

Note: Menu option Generate Response Files creates silent response files for the selected agent; such files contain all the installation parameters required to install the same agent with identical settings on other machines. However, the generated response files apply only to the Common Installer. They do not work with Solution Installer-based agent images.

Running the silent installation from the command-line with parameters

Complete the steps in this section to run the installation from the command-line.

Procedure

1. Start a DOS Command Shell.
2. From the shell, change to the directory containing this installation (where `setup.exe` and `setup.ins` are located).
3. Run the setup as follows. You must specify the parameters in the same order as listed.

```
start /wait setup /z"/sfC:\temp\SILENT_*.TXT" /s /f2"C:\temp\silent_setup.log"
```

where:

/z"/sf"

Specifies the name of the installation driver you customized for your site. This is a required parameter. This file must exist.

SILENT_*.TXT

is the name of the input `silent_server.txt`, `silent_agent.txt`, or `silent_WIA64.txt` file, as described above.

/s Specifies that this is a silent installation. This causes nothing to be displayed during installation.

/f2

Specifies the name of the InstallShield log file. If you do not specify this parameter, the default action is to create `Setup.log` in the same location as the `setup.iss` file. In either case, the Setup program must be able to create and write to this file. If the specified directory does not exist, the Setup program cannot create and write to the file. Therefore, the specified directory path must exist.

Running the silent installation using SMS

Complete the following steps to run the silent installation using SMS.

Procedure

1. Copy all the installation files to a LAN-based disk that SMS will mount on the desired computers. (Copy all files in the directory that contains the setup.exe and setup.ins files.)
2. Replace the original silent_server.txt, silent_agent.txt, or silent_WIA64.txt file on the LAN disk with your modified version.
3. Edit the PDF file located with the setup.exe file, and change the Setup invocation as follows
Setup /z"/sfC:\temp\SILENT_*.TXT" /s /f2"C:\temp\silent_setup.log"

where:

SILENT_*.TXT

is the name of the input silent_server.txt, silent_agent.txt, or silent_WIA64.txt file, as described above.

Installing and configuring the Tivoli Enterprise Monitoring Automation Server in silent mode, by using the silent response file

This configuration option allows for an unattended and automatic installation and configuration of the Tivoli Enterprise Monitoring Automation Server. A locally installed hub monitoring server is a prerequisite for the Tivoli Enterprise Monitoring Automation Server. That means the hub monitoring server is installed together with the Tivoli Enterprise Monitoring Automation Server, unless the hub monitoring server is already installed.

About this task

If you are installing the hub monitoring server as part of this procedure then you must also uncomment the [CMS_CONFIG] section in your silent_server.txt file. If the hub monitoring server is already installed then you do not need to include the [CMS_CONFIG] section in your silent response file.

Complete the following steps to install and configure the Tivoli Enterprise Monitoring Automation Server in silent mode, by using the silent response file:

Procedure

1. On the product installation media, locate the silent_server.txt file in the WINDOWS directory and copy it into a temporary directory on the target system.
2. Edit the silent_server.txt file and perform the following actions so that only the Tivoli Enterprise Monitoring Automation Server is installed:
 - a. uncomment: License Agreement=I agree to use the software only in accordance with the installed license.
 - b. uncomment: Install Directory=C:\IBM\ITM
Accept the default location or specify another directory location.
 - c. uncomment: Install Folder=IBM Tivoli Monitoring
 - d. uncomment: EncryptionKey=IBMTivoliMonitoringEncryptionKey
 - e. uncomment: KMSWICMS=Tivoli Enterprise Monitoring Server
 - f. uncomment: KASWICAS=Tivoli Enterprise Monitoring Automation Server
3. Edit the following section of the silent_server.txt file to configure the Tivoli Enterprise Monitoring Automation Server:
[CAS_CONFIG]
;
; Enable Registry Services access?
; Default value is No. Uncomment your answer:

```

;
;KAS_REGISTRY_SERVICES_ENABLED=Yes

;KAS_REGISTRY_SERVICES_ENABLED=No
;
; The host name of the Registry Services provider registry.
;KAS_REGISTRY_SERVICES_URL_HOSTNAME=
;
; The protocol used by Registry Services provider registry.
; Available protocols must be specified as follows:
;
; =HTTP
; =HTTPS
;
;KAS_REGISTRY_SERVICES_URL_PROTOCOL=HTTP
;
; The port number used by Registry Services provider registry for HTTP connection.
;KAS_REGISTRY_SERVICES_URL_HTTP_PORT=16310
;
; The port number used by Registry Services provider registry for HTTPS connection.
;KAS_REGISTRY_SERVICES_URL_HTTPS_PORT=16311
;
; Set user authentication for Registry Services?
; Default value is No. Uncomment your answer:
;
;KAS_REGISTRY_SERVICES_AUTH=Yes
;KAS_REGISTRY_SERVICES_AUTH=No
;
; The name of a user that is authorized to access Registry Services.
; This must be a valid WebSphere user name.
;KAS_REGISTRY_SERVICES_USER=
;
; The password of the user that is authorized to access Registry Services.
;KAS_REGISTRY_SERVICES_PASSWORD=
;
; The protocol that the Tivoli Enterprise Monitoring Automation Server
; uses for incoming resource query requests. Available protocols
; must be specified as follows:
;
; =HTTP
; =HTTPS
;
;KAS_SERVER_PROTOCOL=HTTP
;
; The port number that the Tivoli Enterprise Monitoring Automation Server
; uses for incoming resource query HTTP requests.
;KAS_SERVER_HTTP_PORT=10001
;
; The port number that the Tivoli Enterprise Monitoring Automation Server
; uses for incoming resource query HTTPS requests.
;KAS_SERVER_HTTPS_PORT=10002

```

4. Invoke the following command to perform the silent installation:

```
start /wait setup /z"/sf<silent_reponse_file>"/s/f2"<silent_install_log>"
```

where:

silent_response_file

Is the absolute path to the edited `silent_server.txt` file.

silent_install_log

Is the absolute path to the silent installation output log file.

5. After the installation is complete you can reconfigure, start, or stop the Tivoli Enterprise Monitoring Automation Server.

Installing and configuring the dashboard data provider in silent mode, by using the silent response file

This configuration option allows for an unattended and automatic installation and configuration of the dashboard data provider.

This is an example of configuration file content for the silent installation of the dashboard data provider:

```
;-----  
;   Dashboard data provider Settings  
;-----  
;  
; Do you want to enable the dashboard data provider service?  
; This parameter is not required. Default value is N.  
; Uncomment your answer:  
;  
;KD8_ITM_REST_SERVICE_ENABLE=N  
;KD8_ITM_REST_SERVICE_ENABLE=Y  
;  
; Enter the Domain override.  
; This parameter is not required and is only valid with the Dashboard  
; data provider is enabled. The value may not exceed 124 characters and  
; may online include 'A-z', 'a-z', '0-9', '-', '_', '.', '@', '#',  
; '$', '*', ':', and ' ' (space) characters.  
;KD8_HUBNAME_OVERRIDE=  
;  
;-----  
;   Authorization Policy Server Settings  
;-----  
;  
; Do you want to enable authorization policies?  
; This parameter is not required. Default value is N.  
; This parameter, as well as the related parameters below, will be  
; ignored if the Dashboard data provider is not enabled.  
; Uncomment your answer:  
;  
;KDQ_GS_ENABLE_POLICY_AUTH=N  
;KDQ_GS_ENABLE_POLICY_AUTH=Y  
;  
; NOTE: The following parameters will be ignored if the  
; authorization policies are not enabled.  
;  
; If authorization policies are enabled, the following  
; parameter is required. Uncomment the line and provide the  
; IP Address or fully qualified hostname of the Policy Server.  
;  
;KDQ_GS_POLICY_SERVER_ADDRESS=  
;  
; Choose the protocol used to connect to the Policy Server.  
; This parameter is not required. The default value is HTTPS.  
; This parameter will be ignored if authorization policies are  
; not enabled. Uncomment your choice:  
;  
;KDQ_GS_WEB_PROTOCOL=HTTPS  
;KDQ_GS_WEB_PROTOCOL=HTTP  
;  
; Choose the port used to connect to the Policy Server.  
; This parameter is not required. The default value is 16311.  
; The valid port values are from 1 to 65535 inclusive.  
;  
;KDQ_GS_POLICY_SERVER_PORT=16311  
;  
; Choose the interval, in minutes, used when obtaining updates  
; from the Policy Server. This parameter is not required. The  
; default is 30 minutes. Valid values are from 5 to 1440 minutes  
; inclusive.  
;  
;
```

```

;KDQ_GS_POLLING_INTERVAL=30
;
; Choose the expiration interval for the local policy store. If
; the policy store cannot be updated from the Policy Server, this
; interval is the amount of time the local policy store will
; continue to be utilized from the last update. These parameters
; are not required. The default is 7 days and 0 hours. The value
; specified for hours must be in the range of 0-23 hours. If
; the expiration interval is set to 0 days and 0 hours, the policy
; store will never expire.
;KDQ_GS_POLICY_STORE_EXP_DAYS=7
;KDQ_GS_POLICY_STORE_EXP_HOURS=0
;
; If authorization policies are enabled, the following
; parameters are required. Uncomment each line and provide the
; User ID and password for the Policy Server.
;
;KDQ_GS_POLICY_SERVER_USERID=
;KDQ_GS_POLICY_SERVER_PWD=
;
; -----
;
; END OF TEPS CONFIGURATION SECTION
;
; -----

```

Performing a silent installation on a Linux or UNIX computer

Just as the interactive installation on Linux and UNIX requires both an installation of code and then a separate configuration, so does the silent installation method. Both the installation and configuration use parameter files to define what you are installing or configuring.

Sample installation and configuration parameter files are included with IBM Tivoli Monitoring and with monitoring agents. The files are located in the following locations:

- Silent installation files:
 - On the product installation media (both base IBM Tivoli Monitoring and agent installation media)
 - After installation, a sample file is located in the *install_dir/samples* directory
- Silent configuration files: After you install the product, a configuration file for each component that requires configuration is located in the *install_dir/samples* directory. There is also a sample configuration file that you can use to configure any component.

Before editing any of the response files, note the following syntax rules:

- Comment lines begin with a pound sign (#).
- Blank lines are ignored.
- Parameter lines are `PARAMETER=value`. Do not use a space before the parameter; you can use a space before or after an equal sign (=).
- Do not use any of the following characters in any parameter value:

\$	dollar sign
=	equal sign
	pipe

Use the following procedures to perform silent installations:

- “Installing components with a response file” on page 849
- “Configuring components with a response file” on page 851

Installing components with a response file

The `silent_install.txt` response file specifies the installation parameters for IBM Tivoli Monitoring components.

About this task

To use this file to perform a silent installation, edit the file to identify what you want to install and then run the following command:

```
./install.sh -q -h install_dir -p response_file [-k]
```

where:

install_dir

identifies the installation location for the IBM Tivoli Monitoring component. The default installation location is `/opt/IBM/ITM`.

Note: You must not specify the path of the directory containing `./install.sh` as your IBM Tivoli Monitoring home directory. On certain platforms, this can cause the plug-in JAR files to overwrite themselves and become zero length files. The installation will fail as a result.

In addition, do not use the number sign (#) or exclamation mark (!) characters anywhere in your installation directory path.

response_file

identifies the response file that you edited to specify installation parameters. Specify the absolute path to this file.

[-k]

is an optional parameter to secure your IBM Tivoli Monitoring installation. If you do not secure the installation at this point you will be asked at the end of this procedure if you want to do so.

[-c]

is an optional parameter to print diagnostic messages to the console.

[-d]

is an optional parameter to indicate the location of the product CD.

[-h]

is an optional parameter to indicate the location of the CandleHome.

[-j]

is an optional parameter to indicate the JRE installed location. This is available on Tandem system only.

[-p]

is an optional parameter for silent mode execution using the specified **PARAMETER_FILE**.

The parameters that you can configure in the `silent_install.txt` file vary depending on the component that you are installing. Each of the files contains comments that explain the options. The following procedure is an example of installing all available components on one computer. Use this to determine the type of information you need to gather when you are setting up your own silent installation.

Use the following steps to perform a silent installation on a UNIX computer:

Procedure

1. Edit the `silent_install.txt` file located in `ITM_build\Windows\Deploy`.
2. Set the following parameters as appropriate for your environment:

Table 154. Silent installation parameters for UNIX

Parameter	Definition
INSTALL_ENCRYPTION_KEY	<p>REQUIRED. The data encryption key used to encrypt data sent between systems. This key must be the same for all components in your IBM Tivoli Monitoring environment.</p> <p>Do not use the following characters in the key:</p> <ul style="list-style-type: none"> & ampersand pipe ' single quote = equal sign \$ dollar sign <p>In addition, do not specify double-byte (DBCS) characters.</p>
INSTALL_FOR_PLATFORM	<p>The operating system for which to install the components. You can specify an architecture code. If you do not specify an architecture code, the operating system for the current computer is used. You can find a list of the architecture codes for the supported architectures in archdsc.tbl in the registry directory; they are also listed in Appendix D, "IBM Tivoli product, platform, and component codes," on page 883.</p>
INSTALL_PRODUCT	<p>The product code for the components (or "products") that you want to install. See Appendix D, "IBM Tivoli product, platform, and component codes," on page 883 for a list of the codes for the base components. You can use the <code>./cinfo</code> command to view the product codes for the applications installed on this computer. You can also find a list of the product codes in the registry directory in proddsc.tbl.</p> <p>You can specify "all" to install all available components.</p> <p>To install multiple components (but not all), repeat this parameter for each component that you want to install. For example:</p> <pre>INSTALL_PRODUCT=ms INSTALL_PRODUCT=cj INSTALL_PRODUCT=cq</pre> <p>This example installs the monitoring server, portal server, and portal desktop client on a Linux computer.</p> <p>In the same pass, you can install support with the components. After you have specified the components that you want to install, specify the support that you want to install using the following parameters:</p> <pre>INSTALL_PRODUCT_TMS (for TEMS support) INSTALL_PRODUCT_TPS (for TEPS support) INSTALL_PRODUCT_TPW (for TEP Browser Client support) INSTALL_PRODUCT_TPD (for TEP Desktop Client support) INSTALL_PRODUCT_TPA (for ITPA Domain support)</pre>
MS_CMS_NAME	<p>If you are installing a monitoring server, use this parameter to specify the name for the monitoring server, such as <code>HUB_hostname</code>. Do not specify an IP address or fully qualified host name.</p>

Table 154. Silent installation parameters for UNIX (continued)

Parameter	Definition
SKIP_SDA_CHECK	You can use this parameter to overwrite product support that was seeded in self-describing mode. By default, this parameter is set to NO and a self-describing mode seeding status check is performed. If this parameter is set to YES , the seeding process at the end of the installation does not check the self-describing mode seeding status.

3. Save and close the file.
4. Run the following command to install IBM Tivoli Monitoring in the /opt/IBM/ITM directory:

```
./install.sh -q -h /opt/ibm/itm -p /tmp/silent_install.txt
```

Configuring components with a response file

You can use the **itmcmd config** command with the **-p filename** parameter to configure IBM Tivoli Monitoring components silently.

About this task

The following sample configuration response files are provided with IBM Tivoli Monitoring:

- **ms_silent_config.txt**: Used to configure the monitoring server
- **cq_silent_config.txt**: Used to configure the portal server
- **cj_silent_config.txt**: Used to configure the portal desktop client
- **silent_config.txt**: A generic configuration file used to configure agents that do not require unique configuration parameters

Note that these sample configuration files are available only after the monitoring server has been installed in a UNIX or Linux system. The files are created in the *itm_installdir/samples* directory.

If an agent requires unique configuration parameters, configure the agent by using the **itmcmd config** command or the Manage Tivoli Enterprise Monitoring Services utility.

Use the following steps to configure a component using the silent method:

Procedure

1. Edit the configuration file for the component that you want to configure.
2. Complete the parameters identified in the file. Each file contains comments that define the available parameters and the values to specify.
3. Save the file and exit.
4. Run one of the following commands.

To configure the monitoring server:

```
./itmcmd config -S -p response_file -t ms_name
```

To configure the portal server, desktop client, or an agent:

```
./itmcmd config -A -p response_file pc
```

where:

response_file

Is the name of the configuration response file. Specify an absolute path to this file.

ms_name

Is the name of the monitoring server that you want to configure.

pc Is the product code for the component or agent that you want to configure. See Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883 for the list of product codes.

Automatically creating agent response files on Linux or UNIX

You can have Tivoli Monitoring create the response files for you after configuring an agent.

About this task

This feature vastly simplifies and speeds up the creation and customization of response files for agent installation and deployment. It also reduces the likelihood of user error when specifying their contents. The resulting response files can be used to either install or deploy similar agents across your environment.

Note: The automatic generation of response files does not apply to multi-instance agents or to server components.

The agent must be successfully installed and configured prior to generating the response file. To have Tivoli Monitoring generate a response file, invoke the `itmcmd` CLI command with the `resp` option:

```
itmcmd resp [-d directory] pc
```

where:

directory

is the name of the directory where you want the generated files stored. The default directory is `itm_installdir/response`.

pc is the product code for the agent whose configuration parameters you want saved. See Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883 for the list of product codes.

Possible errors are that either the *directory* path or the product code, *pc*, is invalid. In either case, an error message is displayed, and the response files are not generated.

When it completes, the `itmcmd resp` command creates these installation and remote-deployment response files:

```
silent_install_pc.txt  
silent_config_pc.txt
```

When response files are available, you can use them to either install or remotely deploy, and then configure, identical agents across your IBM Tivoli Monitoring network:

Procedure

1. Edit the response files, and supply missing security parameters such as passwords and encryption keys.

Note: For security reasons, the IBM Tivoli Monitoring encryption key entered during installation is not saved in the generated response file, and password parameters are blanked out. You must supply these values yourself.

2. Install the agent.
 - If you install it locally on the target computer, perform a silent installation using the generated response file for installation: `silent_install_pc.txt`.
 - If instead you remotely deploy the new agent, use the generated response file for remote deployment: `silent_deploy_pc.txt`.
3. Using the configuration response file (`silent_config_pc.txt`), configure the newly installed agent.

Installing and configuring the Tivoli Enterprise Monitoring Automation Server in silent mode, by using the silent response file

A silent installation runs on its own and does not require you to provide input to dialog boxes or monitor the installation. This configuration option allows for an unattended and automatic installation and configuration of the Tivoli Enterprise Monitoring Automation Server.

About this task

Complete the following steps to install and configure the Tivoli Enterprise Monitoring Automation Server in silent mode, by using the silent response file:

Procedure

1. On the product installation media, locate the `silent_install.txt` file and copy it into a temporary directory on the target system.
2. Edit the `silent_install.txt` file and perform the following actions so that only the Tivoli Enterprise Monitoring Automation Server is installed:
 - a. comment out `INSTALL_PRODUCT=all`
 - b. uncomment `INSTALL_PRODUCT=as`
3. Invoke the following command to perform the silent installation:
`install.sh -q -h <andlehome> -p <silent_response_file>`

Results

The configuration file for the Tivoli Enterprise Monitoring Automation Server is located in `candlehome/samples/as_silent_config.txt`. This is an example of configuration file content for the silent installation of the Tivoli Enterprise Monitoring Automation Server:

```
# Enable Registry Services access
KAS_REGISTRY_SERVICES_ENABLED=YES

# The protocol of the Registry Services
# Available protocols must be specified as follows:
#
#      =HTTP
#      =HTTPS
# Protocol is HTTP or HTTPS.
#
KAS_REGISTRY_SERVICES_URL_PROTOCOL=HTTP

# The hostname of the Registry Services
KAS_REGISTRY_SERVICES_URL_HOSTNAME=

# The HTTP port of the Registry Services
KAS_REGISTRY_SERVICES_URL_HTTP_PORT=16310

# The HTTPS port of the Registry Services
KAS_REGISTRY_SERVICES_URL_HTTPS_PORT=16311

# Set user authentication for Registry Services
KAS_REGISTRY_SERVICES_AUTH=YES

# The name of a user that is authorized to access the Registry Services
# server. This must be a valid WebSphere user name
KAS_REGISTRY_SERVICES_USER=

# The password of the user that is authorized to access the Registry Services server
KAS_REGISTRY_SERVICES_PASSWORD=

# The protocol that the automation server uses
# for incoming resource query requests.
```

```
# Available protocols must be specified as follows:
#
#      =HTTP
#      =HTTPS
#
KAS_SERVER_PROTOCOL=HTTP

# The port number that the automation server uses
# for incoming resource query HTTP requests
KAS_SERVER_HTTP_PORT=10001

# The port number that the automation server uses
# for incoming resource query HTTPS requests
KAS_SERVER_HTTPS_PORT=10002
```

What to do next

After the installation is complete you can reconfigure, start, or stop the Tivoli Enterprise Monitoring Automation Server.

Note: If a monitoring server was also installed by default, you can reconfigure, start, or stop the monitoring server after the installation has completed.

Installing and configuring the dashboard data provider in silent mode, by using the silent response file

This configuration option allows for an unattended and automatic installation and configuration of the dashboard data provider.

This is an example of configuration file content for the silent installation of the dashboard data provider:

```
##### Dashboard data provider #####
# Enable the dashboard data provider?
# Valid answers: Y/N
# This parameter is not required. Default is N.
#KD8_ITM_REST_SERVICE_ENABLE=N

# Enter the Domain override.
# This parameter is not required and is only valid with the Dashboard data provider
# service is enabled. The value may not exceed 124 characters and
# may online include 'A-z', 'a-z', '0-9', '-', '_', '.', '@', '#',
# '$', '*', ':', and ' ' (space) characters.
#KD8_HUBNAME_OVERRIDE=

##### Authorization Policy Server Configuration #####
# Enable authorization policies?
# This parameter is not required. Default value is N.
# This parameter, as well as the related parameters below, will be
# ignored if the Dashboard data provider is not enabled.
# Valid answers: Y/N
#KDQ_GS_ENABLE_POLICY_AUTH=N

# NOTE: The following parameters will be ignored if the
# authorization policies are not enabled.

# If authorization policies are enabled, the following
# parameter is required. Uncomment the line and provide the
# IP Address or fully qualified hostname of the Policy Server.
#KDQ_GS_POLICY_SERVER_ADDRESS=

# Choose the protocol used to connect to the Policy Server.
# This parameter is not required. The default value is HTTPS.
# This parameter will be ignored if authorization policies are
# not enabled. Valid values are HTTPS and HTTP
#KDQ_GS_WEB_PROTOCOL=HTTPS
```

```
# Choose the port used to connect to the Policy Server.
# This parameter is not required. The default value is 16311.
# The valid port values are from 1 to 65535 inclusive.
#KDQ_GS_POLICY_SERVER_PORT=16311

# Choose the interval, in minutes, used when obtaining updates
# from the Policy Server. This parameter is not required. The
# default is 30 minutes. Valid values are from 5 to 1440 minutes
# inclusive.
#KDQ_GS_POLLING_INTERVAL=30

# Choose the expiration interval for the local policy store. If
# the policy store cannot be updated from the Policy Server, this
# interval is the amount of time the local policy store will
# continue to be utilized from the last update. These parameters
# are not required. The default is 7 days and 0 hours. The value
# specified for hours must be in the range of 0-23 hours. If
# the expiration interval is set to 0 days and 0 hours, the policy
# store will never expire.
#KDQ_GS_POLICY_STORE_EXP_DAYS=7
#KDQ_GS_POLICY_STORE_EXP_HOURS=0

# If authorization policies are enabled, the following
# parameters are required. Uncomment each line and provide the
# User ID and password for the Policy Server.
#KDQ_GS_POLICY_SERVER_USERID=
#KDQ_GS_POLICY_SERVER_PWD=
```

Performing a silent installation of software packages using IBM Installation Manager

This section contains information about the prerequisites and procedures for performing a silent installation or update of the following software packages using IBM Installation Manager:

- IBM Infrastructure Management Dashboards for Servers
- Tivoli Authorization Policy Server
- tivcmd Command Line Interface for Authorization Policy

These packages are available on the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image.

A silent installation runs on its own and does not require you to provide input to dialog boxes or monitor the installation. This method of installation is useful for advanced users who supply installation information once through a response file, instead of repeatedly through an installation wizard.

Installing or updating software packages in silent mode

Complete the steps in this section to install or update a software package in silent mode.

Before you begin

Before installing and configuring an IBM Tivoli Monitoring component, install the prerequisite software that is described in “Required software and memory requirements for a dashboard environment” on page 91. Then, create a local repository or network repository using one of the installation approaches described in “Using IBM Installation Manager to install or update IBM Tivoli Monitoring components for a dashboard environment” on page 299.

You must create a response file to use as input to the IBM Installation Manager silent installation commands. You can edit one of the following files available in the *responsefiles* directory on the *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components* DVD or DVD image:

- **Windows** install_ITMAApps4JazzSM_win.xml
- **Linux** **UNIX** install_ITMAApps4JazzSM_unix.xml

For information about the configuration parameters for the Authorization Policy Server and tivcmd CLI, see “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318.

Alternatively, you can use IBM Installation Manager to record a response file. For more information about using IBM Installation Manager to create the response file, see “Recording a response file” on page 858.

Follow the procedure in this section to install or update the packages enabled in the silent response file if the IBM Installation Manager repository is not setup for authentication. If the IBM Installation Manager repository is setup to require authentication, follow the procedure in “Installing or updating silently from authenticated repositories with IBM Installation Manager” on page 858.

Note: If you select the *Configuration* feature for Infrastructure Management Dashboards for Servers, the installer will restart the IBM WebSphere Application Server as part of the installation or update process.

Procedure

1. Open a command prompt and change to the eclipse sub directory in the directory where Installation Manager is installed. The default path for this directory is:
 - **Windows** C:\Program Files\IBM\InstallationManager\eclipse
 - **Linux** **UNIX** /opt/IBM/InstallationManager/eclipse
2. To run the Installation Manager in silent mode, run the following command from the eclipse subdirectory:
 - **Windows** IBMIMc.exe --launcher.ini silent-install.ini -input <response file path and name> -log <log file path and name>
 - **Linux** **UNIX** ./IBMIM --launcher.ini silent-install.ini -input <response file path and name> -log <log file path and name>

Note: Windows does not support using IBMIM.exe for silent installation. Use IBMIMc.exe for silent installation.

For example:

- **Windows** IBMIMc.exe --launcher.ini silent-install.ini -input c:\response_files\install_product.xml -log c:\mylog\silent_install_log.xml
- **Linux** **UNIX** ./IBMIM --launcher.ini silent-install.ini -input /response_files/install_product.xml -log /mylog/silent_install_log.xml

Results

When the silent installation or update is complete, a status of 0 is returned. If the installation or update cannot complete, a non-zero number is returned. When you run the Installation Manager installer, it reads the response file and (optionally) writes to a log file to the directory specified. If you specified a log file and directory, the log file is empty when the operation is successful.

For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<result>
</result>
```

The log file contains an error element if the operation was not completed successfully. A log file for Installation Manager is also available in the following default locations:

- **Windows** (32 bit) as an administrator: C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs

- **Windows** (64 bit) as an administrator: C:\ProgramData\IBM\Installation Manager\logs
- **Windows** as a non-administrator: C:\Documents and Settings\<my id>\Application Data\IBM\Installation Manager\logs
- **Linux** **UNIX** /var/ibm/InstallationManager/logs

What to do next

If you selected the Installation and Configuration features of the IBM Infrastructure Management Dashboards for Servers, perform the following tasks:

1. Verify the dashboard data provider is enabled in the portal server configuration. See “Verifying the dashboard data provider is enabled” on page 427.
2. Then, see *Preparing your dashboard environment* in the *IBM Tivoli Monitoring Administrator's Guide* for additional configuration steps required to set up authentication and authorization for dashboard users. This topic also contains information about creating a data provider connection from IBM Dashboard Application Services Hub to the IBM Tivoli Monitoring dashboard data provider component of the portal server.
3. For more information about launching and using IBM Infrastructure Management Dashboard for Servers, see the *Windows OS Agent User's Guide*, the *Linux OS Agent User's Guide*, or the *UNIX OS Agent User's Guide*.

If you only selected the Installation feature of the IBM Infrastructure Management Dashboards for Servers, review the dashboard application's deployment scripts and then use the **Modify** operation of IBM Installation Manager to deploy the dashboard application into IBM Dashboard Application Services Hub. See “Configuring the IBM Infrastructure Management Dashboards for Servers package after installation” on page 314.

If you installed Tivoli Authorization Policy Server and tivcmd Command Line Interface, perform the following steps:

1. See “Verifying the Tivoli Authorization Policy Server installation” on page 329 for steps to verify that the installation was successful.
2. If you only selected the installation feature of the Tivoli Authorization Policy Server, review the authorization policy server's deployment scripts and then use the **Modify** operation of IBM Installation Manager to deploy the authorization policy server into IBM Dashboard Application Services Hub. See “Configuring the Tivoli Authorization Policy Server feature after installation” on page 328.

Note: For more information about IBM Installation Manager command-line arguments, go to http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp?topic=%2Fcom.ibm.silentinstall12.doc%2Ftopics%2Fr_silent_inst_cmd_arg.html.

3. If you selected both the installation and configuration features of Authorization Policy Server, perform the following tasks:
 - a. If you selected the option for a manual restart of Dashboard Application Services Hub, you must restart the application server for Dashboard Application Services Hub before the Authorization Policy Server can be used with the tivcmd CLI to create authorization policies.
 - b. Use the tivcmd CLI to create authorization policies. The user credentials that you specified during the Authorization Policy Server installation are assigned to the PolicyAdministrator role. You must use these credentials with the tivcmd Command Line Interface to log in to the Authorization Policy Server and assign other administrators permission to create and work with authorization policy roles.

The *IBM Tivoli Monitoring Administrator's Guide* provides examples of creating authorization policies for common scenarios in the *Using role-based authorization policies* chapter. Steps for

configuring a dashboard environment to use authorization policies are in the *Preparing your dashboard environment* chapter. For a complete list of tivcmd CLI commands, see the *IBM Tivoli Monitoring Command Reference*.

Installing or updating silently from authenticated repositories with IBM Installation Manager

You can silently install or update packages from an authenticated repository. However, you must provide the credentials to access protected repositories.

About this task

Complete the following steps to perform a silent installation or update from authenticated repositories with IBM Installation Manager:

Procedure

1. Installation Manager uses the Eclipse infrastructure to save credentials to a key ring file. To generate the key ring file, navigate to the tools folder in the following location:
 - **Windows** C:\Program Files\IBM\InstallationManager\eclipse\tools
 - **Linux** **UNIX** /opt/IBM/InstallationManager/eclipse/tools
2. Run the following command:
 - **Windows** imutilsc.exe saveCredential -url <Installation Repository> -userName <user_id> -userPassword <user_pass> -keyring im.keyring
 - **Linux** **UNIX** imutilsc.sh saveCredential -url <Installation Repository> -userName <user_id> -userPassword <user_pass> -keyring im.keyring
3. After the key ring file is generated, run the silent installation command:
 - **Windows** IBMic.exe --launcher.ini silent-install.ini -input <response file path and name> -keyring <path to the keyring file> -log <log file path and name>
 - **Linux** **UNIX** ./IBMIM --launcher.ini silent-install.ini -input <response file path and name> -keyring <path to the keyring file> -log <log file path and name>

What to do next

See *What to do next* in “Installing or updating software packages in silent mode” on page 855.

Recording a response file

You can use the `-skipInstall <agentDataLocation>` argument to create a response file for an installation or update process without performing the installation.

About this task

When you record a response file with the `-skipInstall <agentDataLocation>` command-line argument, you run the installation process, but the Tivoli Monitoring packages are not installed. The `agentDataLocation` location stores the data for installing the product on other computers. The next time that you use `-skipInstall` with the same `<agentDataLocation>`, you can record a response file to update, modify, roll back, or uninstall the product.

Complete the following steps to record the response file:

Procedure

1. Open a command prompt and change to the eclipse sub directory in the directory where Installation Manager is installed. The default path for this directory is:
 - **Windows** C:\Program Files\IBM\InstallationManager\eclipse

- **Linux** **UNIX** /opt/IBM/InstallationManager/eclipse
2. Run the command to record a response file for the package installation. This command uses the `-skipInstall <agentDataLocation>` argument, which records the installation commands without installing the Tivoli Monitoring packages. Substitute your own file name and location for the response file. Verify that the file paths that you enter exist. Installation Manager does not create directories for the response file:
 - **Windows** IBMIM -record <responseFile> -skipInstall <agentDataLocation>
 - **Linux** **UNIX** ./IBMIM -record <responseFile> -skipInstall <agentDataLocation>

Note: To record a response file and install the IBM product at the same time, remove the argument `-skipInstall <agentDataLocation>`. Use the following command: `IBMIM -record <response file path and name> -log <log file path and name>`.

For example: `IBMIM.exe -record C:\mylog\responsefile.xml -log c:\mylog\record_log.xml`

3. If you copied the repository to a new location, you can set the location of the repository in the repository table. To access the repository table, click **File** → **Preferences** → **Repositories**.
4. Follow the on-screen instructions in the appropriate wizard.
The text (recording) is displayed in the Installation Manager title when you are recording a response file. For more information about the configuration parameters that you are prompted for, see “Installing and configuring the IBM Infrastructure Management Dashboards for Servers” on page 306, “Installing and configuring the Tivoli Authorization Policy Server and tivcmd Command-Line Interface for Authorization Policy” on page 318, and “Installing the tivcmd Command Line Interface for Authorization Policy using the graphical user interface” on page 330.
5. Click **Finish** and close Installation Manager to finish the recording of the response file.

What to do next

To install or update the packages that are selected in the silent response file, perform the procedure in one of the following topics:

- “Installing or updating software packages in silent mode” on page 855
- “Installing or updating silently from authenticated repositories with IBM Installation Manager” on page 858

Agent data location

The `agentDataLocation` value specifies the directory location for the installation folders and files. Although the packages are not installed when recording a silent response file, the files and folders that are required for installation are created and stored in the specified directory. The directory must be writable.

Verify that the file paths in the argument exist. Installation Manager does not create directories for response files.

Attention:

- Do not specify an existing `agentdataLocation` directory for the new `agentdataLocation` location value. For default locations of the `agentdataLocation` directory, go to http://pic.dhe.ibm.com/infocenter/install/v1r6/index.jsp?topic=%2Fcom.ibm.silentinstall12.doc%2Ftopics%2Fr_app_data_loc.html.
- When you record a response file by using the `-skipInstall` argument, you must keep the directory, folders, and files that are created. By keeping the `agentdataLocation` directory, you can modify the installation later. When you record a new response file to modify the installation, you specify the same directory for `agentdataLocation`.

Use a different `agentdataLocation` value for each version of an IBM product that you install.

You must first record an installation of an IBM product by using `-skipInstall` before you record an update.

Appendix C. Firewalls

Review these topics to manage Tivoli Management Services components across network firewalls.

There are four options for achieving interoperability between Tivoli Management Services components across network firewalls:

1. Automatic (do nothing)
2. Configure IP.PIPE with the ephemeral keyword
3. Use a broker partition file
4. Implement a firewall gateway

Use the topics in this section to select and implement the appropriate option for your environment.

- “Determining which option to use”
- “Basic (automatic) implementation” on page 862
- “Implementation with ephemeral pipe” on page 863
- “Implementation with partition files” on page 864
- “Implementation with firewall gateway” on page 868

Determining which option to use

To determine which option is needed in a particular the firewall environment, four factors must be considered.

- “Flow of connection establishment”
- “Permission at the firewall”
- “Server address continuity” on page 862
- “Number of internet zones” on page 862

Flow of connection establishment

Products based on Tivoli Management Services typically establish connections in a traditional client-server flow: connect requests flow from a client in the public network to a server in the trusted network. If the network configuration allows this connection flow, permissions at the firewall are the next consideration in determining how interoperability is achieved.

If the configuration requires that the physical connect requests flow from the secure, trusted server network to the public, untrusted client network, then option 4, a firewall gateway, is required for interoperability. See “Implementation with firewall gateway” on page 868.

Permission at the firewall

If the network configuration allows traditional connection flow, the next consideration is what firewall permissions, if any, are required of the firewall that separates the private, trusted server network from the public, untrusted client network. For simplicity the firewall between these disjoint networks is referred to as the *barrier firewall*.

If all ports are permitted across the barrier firewall, then server address continuity becomes a consideration (see “Server address continuity” on page 862).

If no ports are permitted at the barrier firewall, then to achieve interoperability among components in this firewall environment, full-duplex traffic must be permitted by the firewall administrator for as many

ports as there are servers being concurrently accessed. For example, if Tivoli Enterprise Monitoring Agents are accessing only the Tivoli Enterprise Monitoring Server, then only one port must be permitted (for full-duplex traffic) or opened at the barrier firewall. This is the well-known monitoring server port (the default is 1918 for IP.PIPE, 3660 for IP.SPIPE). If agents are accessing two servers concurrently, a monitoring server and a Warehouse Proxy server, then two ports must be opened at the firewall, one for the monitoring server (typically 1918) and one for the Warehouse Proxy (typically 63358) for interoperability in this firewall environment.

Server address continuity

If the barrier firewall has no restrictions and all ports are permitted, the next factor to consider is server address continuity.

Address continuity refers to the validity and reachability of published IP addresses. Address continuity exists when a published server address is universally reachable by all network clients requesting that service. An example of a server address with address continuity is `update.microsoft.com` (207.46.211.124).

Tivoli Management Services server components register their services and the location of these services (IP address) with a *location broker*. Clients send queries to the location broker to request address information for a service, and receive responses that a list of protocols (address families) and IP addresses at which these services are available. The client then sends a specific server request to one of the addresses in the list received from the location broker. Service registration with the location broker assumes address continuity.

If the published address of the Tivoli service (a remote monitoring server, for example) is identical and reachable for either side of the barrier firewall, then nothing further needs to be done to achieve interoperability in this firewall environment. If the same address cannot be reached from either side of the barrier firewall, then option 2 (ephemeral pipe configuration) or option 3 (broker partition files) is required for interoperability.

Both options are used when traversing a firewall with Network Address Translation (NAT) in effect. While option 2 (ephemeral pipe) is easier to implement, it restricts the endpoint: ephemeral endpoints cannot warehouse data. If warehousing data at the endpoint is required, then partition files must be used for interoperability in this firewall environment (see “Implementation with partition files” on page 864); otherwise ephemeral pipes are sufficient to traverse the translating firewall (see “Implementation with ephemeral pipe” on page 863).

Number of internet zones

The final factor to be taken into consideration is the number of internet zones, that is, how many barrier firewalls must be crossed.

If there are two or more translating firewalls, then option 4, a firewall gateway, must be used for interoperability in this firewall environment (see “Implementation with firewall gateway” on page 868).

Basic (automatic) implementation

IBM Tivoli Monitoring supports most common firewall configurations.

To enable this support, IBM Tivoli Monitoring uses the IP.PIPE socket address family, a TCP-based protocol that opens a single port on the firewall for communication by IBM Tivoli Monitoring components. If your target environment includes a firewall between any IBM Tivoli Monitoring components, you must specify IP.PIPE or IP.SPIPE as your communication protocol during configuration. Typically, no other special configuration is needed unless one of the factors discussed in “Determining which option to use” on page 861 requires it.

Implementation with ephemeral pipe

Configure your communications to use ephemeral pipe when addresses on either side of a barrier firewall are not identical or are not reachable from either side. Use of ephemeral pipe may be dictated by the per-machine agent count, as well as any form of discontinuous addressing such as a NAT firewall.

You configure an IP.PIPE or IP.SPIPE connection as an ephemeral pipe by adding the ephemeral keyword `ephemeral:y` the `KDE_TRANSPORT` environment variable immediately following the protocol keyword (IP.PIPE or IP.SPIPE) in the associated `KPP ENV` file for that process. You must then restart the process for the change to be effective.

Some `KPP ENV` files use `KDC_FAMILIES` instead of `KDE_TRANSPORT`. The process is exactly the same for the `KDC_FAMILIES` environment variable: adding the ephemeral keyword `ephemeral:y` immediately following the protocol keyword (IP.PIPE or IP.SPIPE) that is to be designated ephemeral.

For example, to configure the `KNTAGENT` to make an ephemeral connection to the monitoring server, change `KDE_TRANSPORT` (or `KDC_FAMILIES`) in the file `KNTENV` from

```
KDE_TRANSPORT=IP.PIPE PORT:1918 IP SNA
```

to

```
KDE_TRANSPORT=IP.PIPE ephemeral:y PORT:1918 IP SNA
```

or from

```
KDC_FAMILIES=IP.PIPE PORT:1918 IP SNA
```

to

```
KDC_FAMILIES=IP.PIPE ephemeral:y PORT:1918 IP SNA
```

To configure a remote monitoring server to make an ephemeral connection to the hub, change `KDE_TRANSPORT` (or `KDC_FAMILIES`) in the file `KDSENV` from

```
KDE_TRANSPORT=IP.PIPE PORT:1918 IP SNA
```

to

```
KDE_TRANSPORT=IP.PIPE ephemeral:y PORT:1918 IP SNA
```

or from

```
KDC_FAMILIES=IP.PIPE PORT:1918 IP SNA
```

to

```
KDC_FAMILIES=IP.PIPE ephemeral:y PORT:1918 IP SNA
```

Monitoring agents that configure their connections as ephemeral cannot warehouse data unless `KPX_WAREHOUSE_LOCATION` is also configured at the remote monitoring server to which the monitoring agent reports. The variable `KPX_WAREHOUSE_LOCATION` is an optional list of fully qualified, semicolon-delimited network names that must be added to environment file of the monitoring server to which the agents are connected. This file is located in different places, depending on the platform:

- **Windows** `install_dir\CMS\KBBENV`
- **UNIX** `install_dir/config/kbbenv.ini`

The syntax is:

```
KPX_WAREHOUSE_LOCATION=family_protocol:#network_address[port_number];
```

For example:

```
KPX_WAREHOUSE_LOCATION=ip.pipe:#192.168.0.14[18303];ip:#192.168.0.14[34543];
```

The `KPX_WAREHOUSE_LOCATION` variable does not work if you are using `EPHEMERAL` ports and the Warehouse Proxy Agent is on a server other than the remote monitoring server for that agent. When this is the case, and when "EPHEMERAL:Y" is used in the `KDC_FAMILIES` specification of the agent, the agent will route all communications, including data exports, to its monitoring server. In the agent log, the initial lookup for the Warehouse Proxy does retrieve the correct Warehouse Proxy Agent IP and port from the monitoring server's Local Location Broker. However, the agent switches to using the monitoring server's IP address for the export instead of the Warehouse Proxy Agent IP address due to the `EPHEMERAL:Y` setting.

When the Warehouse Proxy Agent is collocated with the agent's monitoring server, the export works normally because the virtual port that the request gets routed to is the correct listening port for the Warehouse Proxy Agent. The IP address for the Warehouse Proxy Agent and monitoring server are also the same. However, when the Warehouse Proxy Agent is on a different machine to the agent's monitoring server, there is nothing listening on the Warehouse Proxy Agent's port on the monitoring server machine. As a result a Status = 8 error message is received when trying to route the export. There are a few possible solutions:

- Enable and use the KDE Gateway component between the affected agents and their monitoring server.
- Place the Warehouse Proxy Agent on the same monitoring server as the agent.
- If `EPHEMERAL:Y` is only there as a security blanket to ensure monitoring server ↔ agent communications, remove it from the `KDC_FAMILIES` variables of the agents and recycle the agents.

Note: Another possible solution is to change the collection location for all attribute groups from the Tivoli Enterprise Monitoring Agent to the Tivoli Enterprise Monitoring Server. Note that this will considerably increase the monitoring server's processing load and is not usually a preferred method for this reason. However, this method is an option if the previous solutions cannot be implemented.

Related concepts:

"Setting a permanent socket address for a proxy agent" on page 657

In some network environments, use of the Global Location Broker default registration algorithm is not supportable and a specific monitoring server variable must be used.

Implementation with partition files

Address translation is an enhanced security feature of some firewall configurations. With this feature, components that must be reached across the firewall have two unique, but corresponding addresses: the external address (valid for components outside the firewall) and the internal address (valid for components inside the firewall).

In IBM Tivoli Monitoring, the component that typically must be reached for connection is the monitoring server; however, the Warehouse Proxy, which runs on Windows as a server-type application, must also be accessible to clients and also requires an external and internal address. A component on either side of the firewall knows only about the address that is valid for its side, its *partition*.

To accommodate sites with address translation, IBM Tivoli Monitoring uses a partition-naming strategy. This strategy requires two steps:

- The creation of a text file, called a *partition file*, as part of the configuration of a hub or remote monitoring server (or Warehouse Proxy). The partition file contains an entry that defines the address of that component in the other partition.
- The specification of a *partition name* (any alphanumeric string up to 32 characters) as part of the configuration of any agent, a hub or remote monitoring server, or Warehouse Proxy. A partition name must be specified for each component regardless of which side of the firewall it is located on.
- "Sample scenarios" on page 865
- "Windows: Editing the partition file" on page 866

- “UNIX and Linux: Editing the partition file” on page 866
- “Creating the partition file manually” on page 867

Sample scenarios

See scenarios on how to implement partition files in various configurations.

The following scenarios illustrate how to implement partition files in various configurations. In these scenarios, your site has one firewall that contains two partitions, which are named OUTSIDE and INSIDE:

- “Scenario 1: Hub monitoring server INSIDE and monitoring agents OUTSIDE”
- “Scenario 2: Hub and remote monitoring servers INSIDE and monitoring agents OUTSIDE”
- “Scenario 3: Hub monitoring server INSIDE, remote monitoring server and agents OUTSIDE”

Scenario 1: Hub monitoring server INSIDE and monitoring agents OUTSIDE

The hub monitoring server is contained within the firewall in a partition named INSIDE. A partition file named `parthub.txt` is included in this partition and contains the following entry:

```
OUTSIDE ip.pipe: hub's_external_address
```

OUTSIDE is the partition name outside the firewall and *hub's_external_address* is the address of the hub monitoring server that is valid for the agents.

As part of the configuration of each agent, specify the name of the partition that each is located in OUTSIDE.

When an agent starts, `parthub.txt` is searched for an entry that matches the partition name *OUTSIDE* and sees the monitoring server address that is valid for the agents (the external address).

Scenario 2: Hub and remote monitoring servers INSIDE and monitoring agents OUTSIDE

Note: In Scenarios 2 and 3, all agents report to the remote monitoring server.

As part of the configuration of the hub monitoring server, specify the name of the partition that it is located in INSIDE. No partition file is needed because the only component that reports to it (the remote monitoring server) is also inside the firewall.

As part of the configuration of the remote monitoring server, specify the name of the partition that it is located in INSIDE. A partition file, `partremote.txt`, must also be created at the remote monitoring server. It contains the following entries:

```
OUTSIDE ip.pipe: remote's_external_address
```

When configuring the agents (all of which are outside the firewall, reporting to the remote monitoring server), specify the name of the partition that they are located in OUTSIDE. When the agents start, `partremote.txt` is searched for an entry that matches the partition name OUTSIDE and sees the remote monitoring server address that is valid for them (the external address).

Scenario 3: Hub monitoring server INSIDE, remote monitoring server and agents OUTSIDE

As part of the configuration of the hub monitoring server, specify the name of the partition that it is located in INSIDE. Create a partition file, `parthub.txt`, containing the following entry:

```
OUTSIDE ip.pipe: hub's_external_address
```

OUTSIDE is the partition name outside the firewall and *hub's_external_address* is the address of the hub monitoring server that is valid for the remote monitoring server.

As part of the configuration of both the agents and the remote monitoring server, specify the name of the partition they are located in *OUTSIDE*.

A partition file *partremote.txt* also must be created at the remote monitoring server. It contains the following entry:

```
INSIDE ip.pipe: remote's_internal_address
```

If the hub monitoring server needs to communicate with the remote monitoring server (for example, to issue a report request from an agent that is connected to the remote monitoring server), the *partremote.txt* file is searched for an entry that matches the partition name *INSIDE* and sees the remote monitoring server address that is valid for it (the internal address).

Windows: Editing the partition file

You can create or modify the partition file using the Tivoli Enterprise Monitoring Server Configuration option in Manage Tivoli Enterprise Monitoring Services.

About this task

Use the following steps to edit the partition file:

Procedure

1. Right-click the monitoring server you want to configure and select **Reconfigure**. The Tivoli Enterprise Monitoring Server Configuration window is displayed.
2. Select **IP.PIPE** as the communications protocol.
3. Select **Address Translation**.
4. Click **OK**. The Hub (or Remote) TEMS Configuration window is displayed.
5. Click **NAT Settings**. The NAT Settings window is displayed.
6. To create the file:
 - a. For **Partition File**, type the fully qualified name of the partition file, for example `C:\IBM\ITM\CMS\KDCPARTITION.TXT`.
 - b. For **Partition Name**, type the name of the partition to which the file applies.
 - c. Click **Edit File**. A message appears saying that the partition file cannot be found and asking you if you want to create it.
 - d. Click **Yes** to create the file. The file is created and opened in Notepad.
 - e. Create the entry for the partition. The format for the entries is `PARTITION-ID IP.PIPE:nn.nn.nn.nn IP.PIPE:nn.nn.nn.nn`. For example, to create a monitoring server partition for a typical scenario with a monitoring agent outside of a NAT firewall connecting to a monitoring server behind a firewall, use the partition ID of your monitoring agent, two spaces, and then the IP address of the host of the monitoring server. Add additional `IP.PIPE:nn.nn.nn.nn` addresses on a single line for multiple network interface cards. See "Sample partition file" on page 867 for more information on creating entries in the partition file.
 - f. Save and close the file.
7. Click **OK** to save the changes and close the configuration window.

UNIX and Linux: Editing the partition file

You can create or modify the partition file using the Tivoli Enterprise Monitoring Server Configuration option in Manage Tivoli Enterprise Monitoring Services.

About this task

Use the following steps:

Procedure

1. Select the monitoring server you want to configure.
2. Click **Action** → **Configure** → **Basic Settings**.
3. Select **IP.PIPE** as the communications protocol.
4. Select **Use Address Translation**.
5. Enter the full path and file name for the partition file.
6. Click **Create** to create the file (if it does not exist) or **Modify** to edit the file.
7. Enter the partition ID in the first column.
8. Enter the IP address in the second column. If you require a second IP address, enter it in the third column. (If more than two IP addresses are required for a partition ID, use a text editor to add the additional addresses. See “Sample partition file.”)
9. Click **Save** to save the file and exit or **Cancel** to return to the previous screen without modifying the file.

Creating the partition file manually

If your site is using address translation, you must create a partition file. The partition file is a text file containing the name of a partition and its constituent interface address.

You must create or modify this file before implementing firewall support with monitoring servers and agents, the portal server and the hub monitoring server, clients and the portal server, and monitoring agents and Warehouse Proxy Agents.

When Tivoli Management Services components need to communicate across a firewall that performs NAT, those components must be able to retrieve an IP address of the other component that is valid on its side of the firewall. To support this capability, the location broker namespace is logically divided into partitions with unique partition IDs. Partition IDs are specified using the KDC_PARTITION environment variable. The partition file is the means to insert appropriate IP addresses into the location broker namespaces.

When an IBM Tivoli Monitoring component performs a location broker lookup operation, the partition ID of its partition is automatically supplied. The location broker returns only addresses that have been defined for that partition namespace and no other. In effect, the IBM Tivoli Monitoring component sees only addresses that are valid for its partition.

A partition file is a standard text file defined to the system using the KDC_PARTITIONFILE environment variable. Within this file, each line describes a partition name with its constituent IP addresses using space delimited tokens. The format is as follows:

```
PARTITION-ID IP.PIPE:nn.nn.nn.nn IP.PIPE:nn.nn.nn.nn
```

The first token on each line is used as a case-insensitive partition ID. The partition ID can be any alphanumeric string with a maximum length of 32 characters. Subsequent tokens specified are treated as interface addresses in standard NCS format (address-family:address). For communication across firewalls, use only IP.PIPE for address-family.

The expected default location of the file is */install_dir/tables/tems_name*.

Sample partition file

The following sample partition file illustrates the format and content expected.

```
# SAMPLE PARTITION FILE
#
# IMPORTANT: Do not overwrite this file. Copy to another directory
# before making changes.
#
# Lines beginning with a '#' are treated as comments and are ignored.
# Note: Do not specify a line that starts with an '*' as it might prevent
# proper functioning.
#
# Basic Format
# PARTITION-ID IP.PIPE:nn.nn.nn.nn IP.PIPE:nn.nn.nn.nn
#
# Procedure to edit this sample partition file.
# To create a monitoring server partition file for a typical scenario
# (monitoring agent outside of a NAT firewall connecting to a
# monitoring server behind the firewall),do the following:
# 1) Replace the "$OUTSIDE-PID$" with the partition id of your monitoring agent
# 2) Replace the "$OUTSIDE-TEMS-HOST-ADDRESS$" with the ip address of the monitoring
# server host outside of the firewall.
# 3) Add additional IP.PIPE:nn.nn.nn.nn addresses on a single line for multiple
# Network Interface Cards (NICs) as in the format above. Separate entries with
# two spaces.Lines can be continued by placing a backslash ('\') char at
# the end of the line.
#
#####
$OUTSIDE-PID$ IP.PIPE:$OUTSIDE-CMS-HOST-ADDRESS$
```

Implementation with firewall gateway

The firewall gateway feature enables additional end-to-end connectivity options for use in environments with specific TCP/IP connection management policies.

The firewall gateway is capable of negotiating multiple firewall hops and network address translation (NAT). It also allows you to configure network traffic so that it is initiated from the more secure network zone.

The Firewall Gateway provides the following functionality:

- Gateway instances interoperate over a single physical relay connection. Logical connections are multiplexed over the relay. The origination direction of the relay connection is configurable to match enterprise firewall transit requirements.
- Relay support enables a logical connection to span multiple firewall zones. Each relay instance can optionally provide access to the upstream management network. Multiple relays can be chained to provide seamless hops across multiple zones.
- Proxy support provides a transparent interface to IBM Tivoli Monitoring components. Server proxy components reside downstream and listen for inbound connections. Client proxy components reside upstream and make connections to services on behalf of downstream endpoints.
- All ports used by gateway instances are configurable. Port pooling is available to constrain client proxy connections to designated port values.
- Multiple failover addresses can be configured for all gateway connections.

NAT alone is not a reason to use the firewall gateway, which is content-neutral and can proxy any TCP connection. In most cases, NAT processing is handled by the PIPE protocol (IP.PIPE or IPSPIPE), which can be used without the firewall gateway. Use the gateway when you have any of the following scenarios:

- A single TCP connection cannot be made to span between IBM Tivoli Monitoring components. An example would be that there are multiple firewalls between these components and a policy that does not allow a single connection to traverse multiple firewalls.

- Connection requirements do not allow the IBM Tivoli Monitoring default pattern of connections to the hub monitoring server. An example here would be agents residing in a less-secure zone connecting to a monitoring server residing in a more-secure zone. Security policy would only allow a connection to be established from a more-secure zone to a less-secure zone, but not the other way round.
- You must reduce open firewall ports to a single port or connection. For example, rather than opening the port for every system being monitored, you would like to consolidate the ports to a single “concentrator”. Connection requirements do not allow the IBM Tivoli Monitoring default pattern of connections to the hub monitoring server.
- You must reduce open firewall ports to a single port or connection. You must manage agent failover and monitoring server assignment symbolically at the hub monitoring server end of the gateway. Because gateway connections are made between matching service names, an administrator can change the failover and monitoring server assignment of downstream gateway agents by changing the client proxy bindings next to the hub monitoring server.

In the context of firewalls, the server and client relationship can best be described in terms of *upstream* and *downstream*. Those entities that open a socket to listen for requests are at the upstream or server end. Those entities connecting to the server are at the downstream or client end. Using one or more relay configurations, logical connection requests flow from a listening downstream server proxy interface, and terminate in an outbound connection from an upstream client proxy interface to a listening server. Intermediate relay configurations consist of an upstream relay interface containing at least one downstream relay interface.

Configuration

The gateway component is configured through an XML document that specifies a set of zones, each of which contain at least one upstream interface with one or more imbedded downstream interfaces.

The following sections provide information on the activation, structure, and content of the XML document:

- “Activation”
- “IPv4 Address Data” on page 870
- “IPv6 Address Data” on page 870
- “XML Document Structure” on page 870

Activation

The gateway feature can be activated within any IBM Tivoli Monitoring process. However, use must be limited to the host computer operating system agent to prevent potential resource consumption conflicts with Tivoli Enterprise Monitoring Server and Tivoli Enterprise Portal Server processes.

The configuration variable `KDE_GATEWAY` is set to the XML configuration file name. A line of the form `KDE_GATEWAY=filename` must be added to the following configuration files, depending on your environment:

- **Windows** configuration variables for the Windows operating system agent are located in the `ITMinstall_dir\tmaitm6\KNTENV` file.
- **UNIX** configuration variables for the UNIX operating system agent are located in the `ITMinstall_dir/config/ux.ini` and `ITMinstall_dir/config/ux.config` files. Add the entry to both files for reliable results.
- **Linux** configuration variables for the Linux operating system agent are located in the `ITMinstall_dir/config/lz.ini` and `ITMinstall_dir/config/lz.config` files. Add the entry to both files for reliable results.

After you make these changes, stop and restart the monitoring agents.

For information on configuring firewall support on z/OS, see *IBM Tivoli OMEGAMON XE and Tivoli Management Services on z/OS Common Planning and Configuration* (<http://pic.dhe.ibm.com/infocenter/>)

tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/zcommonconfig/zcommonconfig.htm) and *Configuring the Tivoli Enterprise Monitoring Server on z/OS* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm).

IPv4 Address Data

Internet Protocol Version 4 (IPv4) addresses supplied as data to **<bind>** and **<connection>** tags can be in absolute dotted decimal or symbolic form.

An address-specific port number override can be specified following a trailing colon (":") character.

IPv6 Address Data

Internet Protocol Version 6 (IPv6) addresses supplied as data to **<bind>** and **<connection>** tags can be in absolute uncompressed hexadecimal, absolute compressed hexadecimal, or symbolic form.

Absolute hexadecimal expressions must be enclosed within parentheses "(" and ")") with one-to four-digit groups separated by a colon (":"). Compression of a run of 0 digits can occur at most once and is indicated by double colons ("::"). An address-specific port number override can be specified following a trailing colon; this specification is outside the parentheses that wrap an absolute address.

XML Document Structure

This section illustrates the structure of firewall gateway XML configuration.

The relationships between XML configuration elements are illustrated in Figure 161. Attributes are described on affected elements; default values for most attributes can be supplied on outer elements with noted exceptions.

```
<tepgwml:gateway xmlns:tepgwml="http://xml.schemas.ibm.com/tivoli/tep/kde/">
  <zone>
    <interface>      upstream interface
    <bind>
      <connection>
      </connection>
    </bind>
    <interface>      downstream interface
    <bind>
      <connection>
      </connection>
    </bind>
  </interface>
</zone>
<portpool>
</portpool>
</tepgwml:gateway>
```

Figure 161. Structure of firewall gateway XML configuration document

<gateway>

A **<gateway>** element in the assigned namespace <http://xml.schemas.ibm.com/tivoli/tep/kde/> contains configuration elements described within this document. The gateway XML processor semantically ignores valid XML until the container is opened, allowing for configuration documents to be imbedded in other documents. This element cannot contain data.

name

The **name** attribute is required, cannot contain imbedded delimiters, and must begin with a nonnumeric. This attribute is used to identify a specific gateway instance. This attribute cannot be inherited from an outer element.

threads

The **threads** attribute specifies the number of worker threads in a general purpose thread pool. The specification must satisfy $1 \leq \text{value} \leq 256$, and defaults to 32. Threads in this

pool are shared by all defined zones, and are used only by interface startup logic, and to recover from outbound buffer exhaustion conditions. The default value is generally more than adequate.

<zone>

A zone is a container of interfaces sharing communication resources. This element cannot contain data.

name

The **name** attribute is required, cannot contain imbedded delimiters, and must begin with a nonnumeric. This attribute is used to identify a specific zone instance. This attribute cannot be inherited from an outer element.

maxconn

The **maxconn** attribute imposes an upper limit on the number of concurrent gateway connections within the zone. Each proxy physical connection and each logical connection crossing a relay interface consume this value. The specification must satisfy $8 \leq \text{value} \leq 4096$, and defaults to 256.

bufsize

The **bufsize** attribute sets the data buffer size within the zone. The specification must satisfy $256 \leq \text{value} \leq 16384$, and defaults to 2048.

minbufs

The **minbufs** attribute sets the minimum number of buffers in the zone buffer pool that are reserved for inbound traffic. The specification must satisfy $4 \leq \text{value} \leq 1024$, and defaults to 64.

maxbufs

The **maxbufs** attribute sets the maximum number of buffers in the zone buffer pool that are reserved for inbound traffic. The specification must satisfy $\text{minbufs} \leq \text{value} \leq 2048$, and defaults to 128.

<interface>

An interface describes a set of network bindings that exhibit a fixed behavior according to a specified role, and based on whether it is defined as upstream, which means that the enclosing element is <zone>, or downstream, where the enclosing element is <interface>. In all roles, logical connections arrive through one or more downstream interfaces and are forwarded through the upstream interface. After a logical connection has been established end to end, data flow is full duplex. A valid configuration requires an upstream interface to contain at least one downstream interface. This element cannot contain data.

name

The **name** attribute is required, cannot contain imbedded delimiters, and must begin with a nonnumeric. This attribute is used to identify a specific interface instance. This attribute cannot be inherited from an outer element.

role

The **role** attribute is required, and describes the behavior of network bindings contained within. The role attribute must be specified as "proxy", "listen", or "connect". Downstream proxy interfaces represent local listening endpoints, and function as a server proxy. Upstream proxy interfaces represent local connecting endpoints, and function as a client proxy. Relay interfaces are assigned either "listen" or "connect". No configuration restriction is made on the relay connection role other than peer relay connections must specify the opposite role. Relay connections are considered persistent, are initiated at gateway startup, and automatically restarted in the event of a network disruption.

<bind>

A <bind> element represents connection resources on one or more local interfaces. When specified within interfaces that "listen" (downstream proxy, relay listen), bind elements represent listening ports on local interfaces. For "connect" interfaces (upstream proxy, relay connect), they

represent the local binding to be used for the outbound connection. Specific local interface addresses can be supplied as data; the default interface is any.

localport

The **localport** attribute is required within “listen” interfaces, and is optional within “connect” interfaces. The value supplied can be either a number that satisfies $1 \leq \text{value} \leq 65535$, or for “connect” based roles, can only contain the name of a portpool element defined within the gateway.

ipversion

The **ipversion** attribute declares the address family to be used for activity within the tag scope. Valid values are 4 or 6, with a default of 4.

ssl

The **ssl** attribute controls SSL negotiation for connections within the scope of this binding. When specified as yes, a successful negotiation is required before a connection is allowed on the gateway. The default value is no, meaning no SSL negotiation occurs on behalf of the gateway connection. Note that this does not restrict the conveyance of SSL streams across a gateway, only whether or not the gateway acts as one end of the SSL negotiation. When this operand is specified on a relay binding, it can be used to secure relay traffic, and must be specified on both ends of the relay connection.

service

The **service** attribute is a character string used to represent a logical connection between client and server proxy interfaces. Each connection accepted by a server proxy must find an upstream client proxy connection with a matching service string. No value restrictions are imposed.

<connection>

The **<connection>** tag is used to supply remote network interfaces as data. When applied to a “listen” mode binding, the connection tag represents the list of remote interface addresses that are allowed to make a connection, and is optional. This tag is required for “connect” mode bindings, and describes the remote end of the connection. Multiple addresses can be supplied for failover purposes.

remoteport

The **remoteport** attribute supplies the default port number of remote interfaces described within this tag. The value supplied must satisfy $1 \leq \text{value} \leq 65535$.

<portpool>

The **<portpool>** tag is used to create a list of local port numbers to be used for outbound connections. Port numbers are supplied as data, and can be specified discretely or as a range expression separated by hyphen (“-”). Range expressions are limited to 1024 bytes to prevent syntax errors from resulting in larger ranges than expected. Multiple specifications of either form are allowed.

name

The **name** attribute is required, cannot contain imbedded delimiters, and must begin with a nonnumeric. This attribute is used to identify a specific portpool instance. This attribute cannot be inherited from an outer element, and is referenced by a *localport* attribute on a bind element.

Warehouse Proxy Configuration

Append configuration text to the Warehouse Proxy KDC_FAMILIES configuration variable.

To ensure that the Warehouse Proxy Agent listens at a fixed port number across the monitoring enterprise, append the following configuration text to the KDC_FAMILIES configuration variable for the Warehouse Proxy:

```
IP.PIPE SKIP:15 COUNT:1
```

The effect of this configuration change is to force the Warehouse Proxy to listen at the Tivoli Enterprise Monitoring Server well-known port number (default 1918) plus the quantity 4096 multiplied by 15. For example purposes, if the monitoring server port is defaulted to 1918, this causes the Warehouse Proxy to listen at 63358. The following examples assume this recommendation has been implemented.

Example gateway configuration scenario

This section illustrates firewall gateway configuration in a three-hop network configuration.

This example uses a three-hop firewall scenario, shown in Figure 162. This scenario makes the following assumptions:

- Connections can only cross a firewall from the more trusted side to the less trusted side.
- Relay data crossing a zone enters and leaves on separate ports.

The effects of NAT on cross-zone addresses are not shown for clarity. NAT connections are fully supported. Dynamic NAT connections may require that inbound connection verification be removed. This is accomplished by removing the **<connection>** tag under the "listening" **<bind>**.

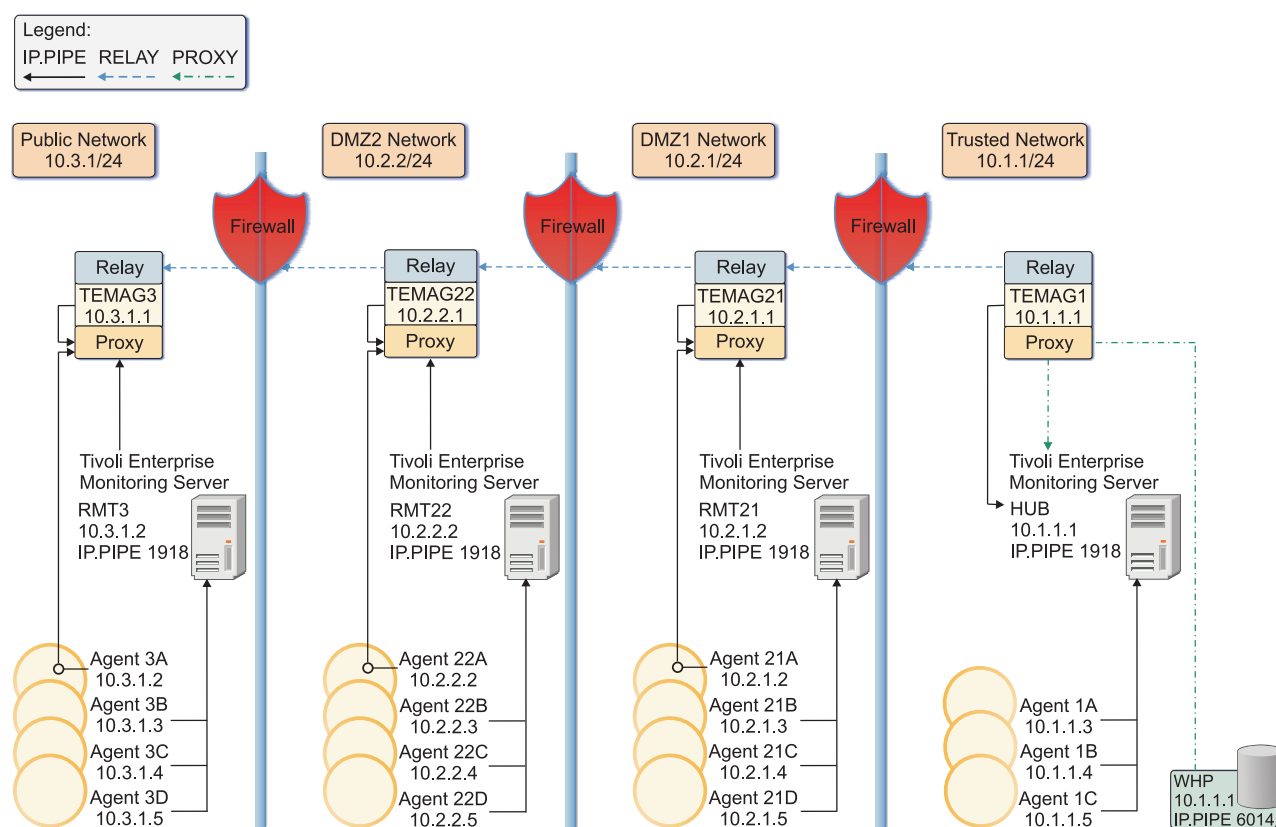


Figure 162. Three-hop firewall scenario

Public Network (TEMAG3)

The public network has the following characteristics:

- Gateway service is configured as part of operating system agent TEMAG3 on 10.3.1.1.
- TEMAG3 accepts a relay connection on port 10030 only from TEMAG22, port 10030.
- The Tivoli Monitoring components within this zone will contact the hub monitoring server and Warehouse Proxy server proxy ports 1918 and 63358 through the TEMAG3 interface address.

The agents and the monitoring server at these zones needs to be configured as follows: Agents 3B, 3C, and 3D pointed directly to RMT3. RMT3 needs to be configured to point to TEMAG3, even if the configuration dialog asks for the hostname or address of the primary hub monitoring server. Agent3A as well as TEMAG3 itself should both point to TEMAG3.

In general, a gateway agent should point to itself, except for the final gateway, which should point to a monitoring server as usual. In this example, TEMAG1 should point to HUB.

In terms of node topology, all the agents and monitoring servers in this example that are pointed to the gateway agents will appear as if they are directly connected to the hub monitoring server.

- A remote monitoring server resides on a computer other than TEMAG3 to prevent conflict on port 1918.

The TEMAG3 gateway has the following configuration:

```
<tep:gateway name="temag3"
xmlns:tep="http://xml.schemas.ibm.com/tivoli/tep/kde/" >
<zone name="least_trusted">
  <!--
  upstream interface, listens for incoming relay
  connections, accepts traffic from downstream interfaces.
  -->
  <interface name="uprelay" ipversion="4" role="listen">
    <bind localport="10030">10.3.1.1
      <connection remoteport="10030">10.2.2.1
    </connection>
    </bind>
  </interface>
  <!--
  downstream interface, listens for incoming proxy
  connections, routes traffic over upstream relay.
  -->
  <interface name="serverproxy" ipversion="4" role="proxy">
    <bind localport="1918" service="tems"/>
    <bind localport="63358" service="whp"/>
  </interface>
</interface>
</zone>
</tep:gateway>
```

DMZ2 Network (TEMAG22)

The DMZ2 network has the following characteristics:

- Gateway service configured as part of OS agent TEMAG22 on 10.2.2.1.
- TEMAG22 originates a relay connection to TEMAG3 port 10030 using local port 10030.
- TEMAG22 accepts a relay connection on port 10022 only from TEMAG21, port 10022.
- Tivoli Monitoring components within this zone will contact the hub monitoring server and Warehouse Proxy server proxy ports 1918 and 63358 through the TEMAG22 interface address.
- A remote monitoring server resides on a computer other than TEMAG22 to prevent conflicts on port 1918.

The TEMAG22 gateway has the following configuration:

```
<tep:gateway name="temag22"
xmlns:tep="http://xml.schemas.ibm.com/tivoli/tep/kde/" >
<zone name="dmz2">
  <!--
  upstream interface, listens for incoming relay
  connections, accepts traffic from downstream
  interfaces
  -->
  <interface name="uprelay" ipversion="4" role="listen">
    <bind localport="10022">10.2.2.1
      <connection remoteport="10022">10.2.1.1</connection>
```

```

</bind>
<!--
downstream interface, originates relay connection to
downstream relay, routes traffic over upstream relay.
-->
<interface name="downrelay" ipversion="4" role="connect">
  <bind localport="10030">10.2.2.1
    <connection remoteport="10030">10.3.1.1</connection>
  </bind>
</interface>
<!--
downstream interface, listens for incoming proxy
connections, routes traffic over upstream relay.
-->
<interface name="serverproxy" ipversion="4" role="proxy">
  <bind localport="1918" service="tems"/>
  <bind localport="63358" service="whp"/>
</interface>
</interface>
</zone>
</tep:gateway>

```

DMZ1 Network (TEMAG21)

The DMZ1 network has the following characteristics:

- Gateway service is configured as part of OS agent TEMAG21 on 10.2.1.1.
- TEMAG21 originates a relay connection to TEMAG22 port 10022 using local port 10022.
- TEMAG21 accepts a relay connection on port 10021 only from TEMAG1 port 10021.
- Tivoli Monitoring components within this zone will contact the hub monitoring server and Warehouse Proxy server proxy ports 1918 and 63358 via the TEMAG21 interface address.
- A remote monitoring server resides on a computer other than TEMAG21 to prevent conflicts on port 1918.

The TEMAG21 gateway has the following configuration:

```

<tep:gateway name="temag21"
xmlns:tep="http://xml.schemas.ibm.com/tivoli/tep/kde/" >
  <zone name="dmz1">
    <interface name="uprelay" ipversion="4" role="listen">
      <bind localport="10021">10.2.1.1
        <connection remoteport="10021">10.1.1.1</connection>
      </bind>
    <interface name="downrelay" ipversion="4" role="connect">
      <bind localport="10022">10.2.1.1
        <connection remoteport="10022">10.2.2.1</connection>
      </bind>
    </interface>
    <interface name="serverproxy" ipversion="4" role="proxy">
      <bind localport="1918" service="tems"/>
      <bind localport="63358" service="whp"/>
    </interface>
  </interface>
</zone>
</tep:gateway>

```

Trusted Network (TEMAG1)

The Trusted Network has the following characteristics:

- Gateway service is configured as part of operating system agent TEMAG1 on 10.1.1.1.
- TEMAG1 originates a relay connection to TEMAG21 port 10021 using local port 10021.

- TEMAG1 makes client proxy connections to the hub monitoring server using source ports in the range 20000-20099.
- TEMAG1 makes client proxy connections to the Warehouse Proxy Agent at destination port 63358 using source ports in the range 20100-20199.

The TEMAG1 gateway has the following configuration:

```
<tep:gateway name="temag1"
xmlns:tep="http://xml.schemas.ibm.com/tivoli/tep/kde/" >
  <zone name="most_trusted">
    <!--
      upstream interface, traffic from downstream
      interfaces and originates proxy connections on behalf
      of downstream server proxy clients
    -->
    <interface name="clientproxy" ipversion="4" role="proxy">
      <bind localport="poolhub" service="tems">
        <connection remoteport="1918">10.1.1.1</connection>
      </bind>
      <bind localport="poolwhp" service="whp">
        <connection remoteport="63358">10.1.1.1</connection>
      </bind>
    <!--
      downstream interface, originates connection to
      downstream relay, routes traffic to upstream proxy
    -->
    <interface name="downrelay" ipversion="4" role="connect">
      <bind localport="10021">10.1.1.1
        <connection remoteport="10021">10.2.1.1</connection>
      </bind>
    </interface>
  </interface>
</zone>
<portpool name="poolhub">20000-20099</portpool>
<portpool name="poolwhp">20100-20199</portpool>
</tep:gateway>
```

Firewall network address translation (NAT) or multiple network interface cards

The URL for starting Tivoli Enterprise Portal in browser mode includes the Tivoli Enterprise Portal Server host name or IP address. The address for starting Tivoli Enterprise Portal is set for the desktop client during installation or through Manage Tivoli Enterprise Monitoring Services.

If any of the following is true in your configuration, you need to define a Tivoli Enterprise Portal Server interface:

- A firewall with Network Address Translation (NAT) is used between the client and the Tivoli Enterprise Portal Server.
- The host of the Tivoli Enterprise Portal Server has more than one Network Interface Card (NIC).

On Windows you define an interface using Manage Tivoli Enterprise Monitoring Services. On Linux and AIX, you edit the `cq.ini` file manually.

Defining a Tivoli Enterprise Portal Server interface on Windows

Use the steps outlined in this section to define a Tivoli Enterprise Portal Server interface on Windows.

Procedure

1. On the computer where the Tivoli Enterprise Portal Server is installed, click **Start** → **Programs** → **IBM Tivoli Monitoring** → **Manage Tivoli Monitoring Services**.
2. Right-click **Tivoli Enterprise Portal Server**.

3. Click **Advanced** → **Configure TEPS Interfaces**. Initially, the list has one definition named "cnps," using port 15001 for the Tivoli Enterprise Portal Server and the IBM Tivoli integrated Web server at `http://mysystem:15200` (where the variable *mysystem* is the host name). Port 80, for an external Web server, is assumed if the URL does not specify 15200 for the integrated Web server.

Note: If IIS is being used as the external Web server for Tivoli Enterprise Portal over a firewall, you may experience performance problems like slow download times. Socket pooling may also cause problems under certain conditions. If you encounter problems, try using a port other than the default port 80 on the IIS server.

4. Click **Add**.
5. Define the interface. Complete the following fields:

Interface Name

Type a one-word title for the interface.

Host If you are defining an interface for a specific NIC or different IP address on this computer, type the TCP/IP host address. Otherwise, leave this field blank.

Proxy Host

If you are using address translation (NAT), type the TCP/IP address used outside the firewall. This is the NATed address.

Port Type a new port number for the Tivoli Enterprise Portal Server. The default 15001 is for the server host address, so a second host IP address or a NATed address requires a different port number.

Proxy Port

If the port outside the firewall will be translated to something different than what is specified for **Port**, set that value here.

Enable SSL for TEP clients

Enable secure communications between clients and the portal server.

6. Click **OK** to add the new Tivoli Enterprise Portal Server interface definition to the list.

Defining a Tivoli Enterprise Portal Server interface on Linux or UNIX

To define an additional Tivoli Enterprise Portal interface on Linux or UNIX, edit the *install_dir/config/cq.ini* file as described in this section.

Procedure

1. Locate the `KFW_INTERFACES=` variable and add the one-word name of the new interface, separating it from the preceding name by a space. For example:

```
KFW_INTERFACES=cnps myinterface
```

2. Following the entries for the default cnps interface, add the following variables as needed, specifying the appropriate values:

KFW_INTERFACE_interface_name_HOST=

If you are defining an interface for a specific NIC or different IP address on this computer, specify the TCP/IP host address.

KFW_INTERFACE_interface_name_PORT=

Type a port number for the Tivoli Enterprise Portal Server. The default 15001 is for the server host address, so a second host IP address or a NATed address requires a different port number.

KFW_INTERFACE_interface_name_PROXY_HOST=

If you are using address translation (NAT), type the TCP/IP address used outside the firewall. This is the NATed address.

KFW_INTERFACE_interface_name_PROXY_PORT=

If the port outside the firewall will be translated to something different than what is specified for **Port**, set that value here.

KFW_INTERFACE_interface_name_SSL=Y

If you want clients to use Secure Sockets Layers (SSL) to communicate with the Tivoli Enterprise Portal Server, add this variable.

Firewall scenarios for Tivoli Enterprise Portal

The diagrams contained in this section illustrate several firewall scenarios using various combinations of the IBM Tivoli integral Web server, an external Web server (such as Apache or IBM HTTP Server), NAT, and a second NIC on the Tivoli Enterprise Portal Server computer.

These scenarios can help you to define the Tivoli Enterprise Portal Server interface.

Figure 163 shows a scenario with the following configuration:

- Has an intranet firewall
- Has no NAT
- Uses the integral Web server




Figure 163. Intranet with integral Web server

The default Tivoli Enterprise Portal Server interface "cnps" is used. No additional interface definitions are needed. Browser mode users, whether going through the firewall or not, start Tivoli Enterprise Portal at:

`http://10.10.10.10:15200`

or substitute the host name for the IP address.

For configurations using the integrated Web server and these port numbers, use the default cnps interface definition.

In this scenario, the monitoring server and agents can be installed on the Tivoli Enterprise Portal Server computer.

Figure 164 shows a scenario that has the following configuration:

- Has an intranet firewall
- Has no NAT
- Uses an external Web server (such as IBM HTTP Server, Apache or IIS)




Figure 164. Intranet with external Web server

Browser mode users, whether going through the firewall or not, start Tivoli Enterprise Portal Server with `http://10.10.10.10` or `http://10.10.10.10/mydirectory`

(where `mydirectory` is the alias), or substitute the host name for the IP address.

For intranet configurations using an external Web server, with no NAT, you do not need to add a new interface definition. Web server port 80 is used automatically when none is specified in the URL.

In this scenario, the monitoring server and agents can be installed on the Tivoli Enterprise Portal Server computer.

Figure 165 on page 880 shows the following two-part configuration:

- Intranet firewall without NAT and using the integral Web server
- Internet firewall with NAT and using an external Web server

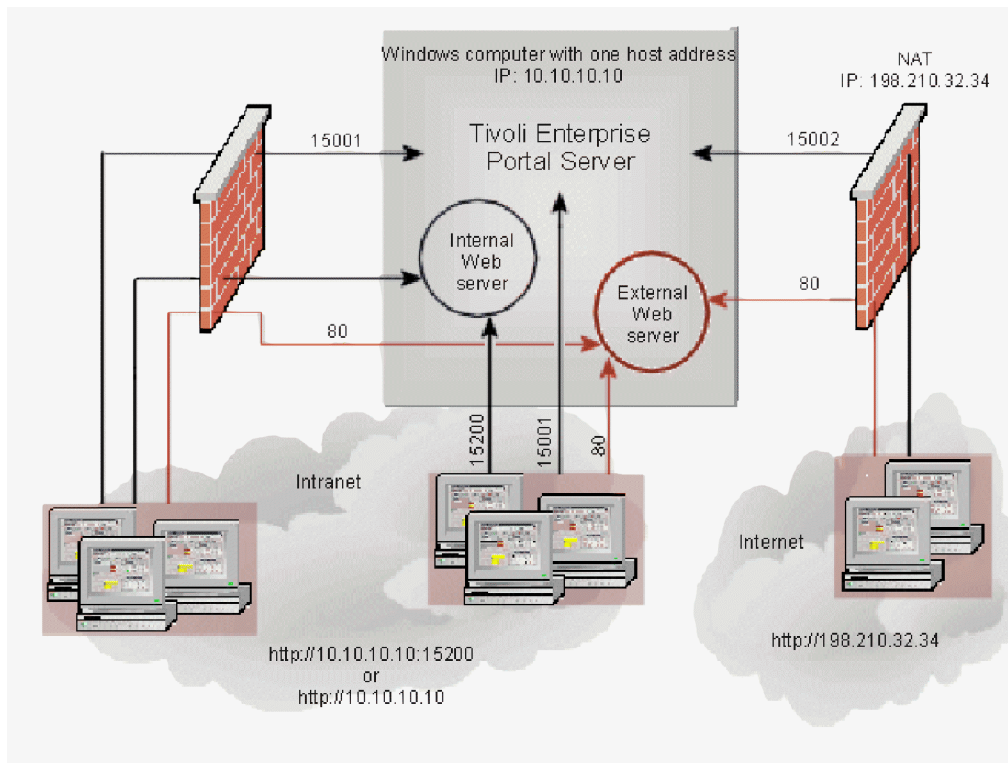


Figure 165. Intranet with integral Web server; Internet with external Web server

Intranet users can enter the URL for either the integral Web server or the external Web server:

`http://10.10.10.10:15200` or `http://10.10.10.10`

Internet users enter the URL for the NATed address:

`http://198.210.32.34/?ior=internet.ior`

(or substitute the host name for the IP address).

The Internet configuration requires a new Tivoli Enterprise Portal Server interface named "internet", with proxy host address 198.210.32.34 and port number 15002. The intranet firewall uses the "cnps" definition.

In this scenario, the monitoring server and agents cannot be installed on the Tivoli Enterprise Portal Server computer.

Figure 166 on page 881 shows the following three-part configuration:

- Intranet firewall with NAT through the firewall to the external Web server
`http://192.168.1.100/?ior=intranet.ior`
- Without NAT inside the DMZ to the integral Web server
`http://10.10.10.10:15200`
- Internet firewall with NAT through the firewall to the external Web server
`http://198.210.32.34/?ior=internet.ior`

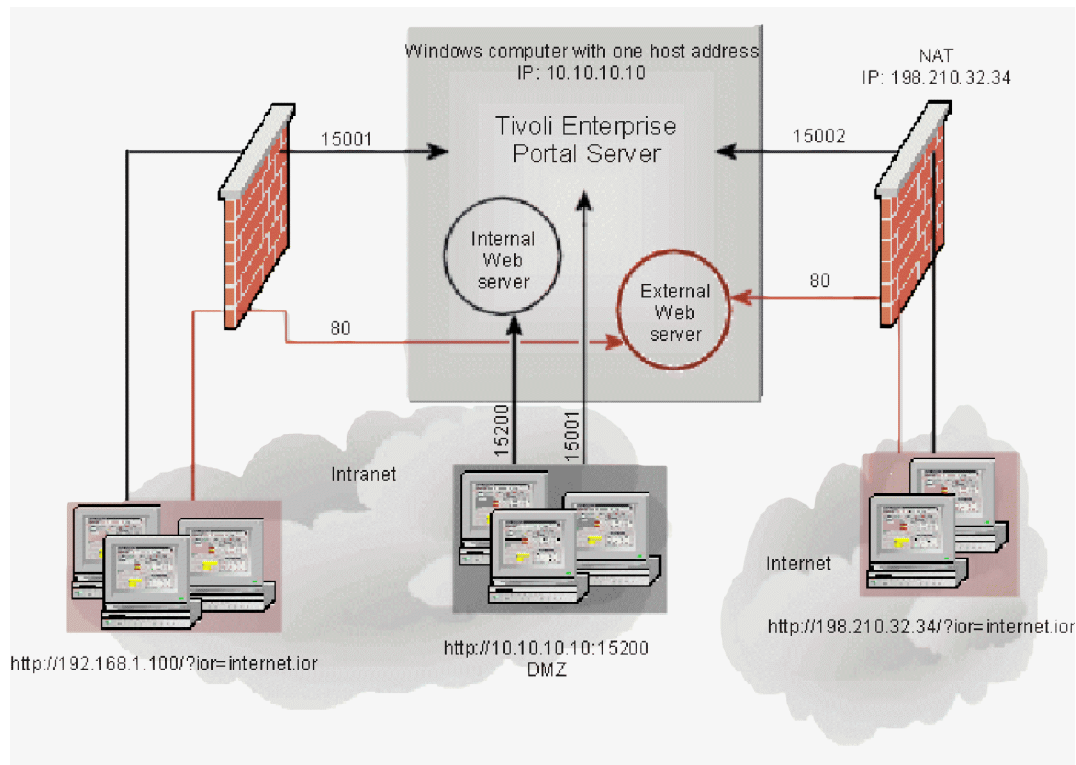


Figure 166. Intranet and Internet with integral and external Web servers

The intranet firewall configuration requires a new Tivoli Enterprise Portal Server interface named "intranet", which uses proxy host 192.168.1.100 and port 15003.

The Internet DMZ configuration requires a new Tivoli Enterprise Portal Server interface definition.

The Internet configuration uses the same Tivoli Enterprise Portal Server "internet" interface definition as the previous scenario: proxy host 198.210.32.34 and port 15002.

In this scenario, the monitoring server and agents cannot be installed on the Tivoli Enterprise Portal Server computer.

Figure 167 on page 882 shows the following two-part configuration:

- Intranet firewall with NAT through the firewall to the external Web server using <http://192.168.1.100>, and without NAT inside the DMZ to the integral Web server uses <http://10.10.10.10:15200>
- Internet firewall with NAT through the firewall to the external Web server using <http://198.210.32.34>.




Figure 167. Two host addresses, intranet and Internet, with integral and external Web servers

The intranet firewall configuration uses the same Tivoli Enterprise Portal Server interface definition (named "intranet") as in the scenario shown in Figure 166 on page 881: `http://10.10.10.10`; proxy host 192.168.1.100; and port 15003.

The intranet DMZ configuration uses the default Tivoli Enterprise Portal Server interface definition: host 192.168.33.33; proxy host 198.210.32.34; port 15002; and proxy port 444.

In this scenario, the monitoring server and agents cannot be installed on the Tivoli Enterprise Portal Server computer.

Appendix D. IBM Tivoli product, platform, and component codes

This appendix contains the IBM Tivoli Monitoring product, platform, and component codes.

Component product codes for infrastructure components and base monitoring agents

Table 155 lists the product codes that identify the different IBM Tivoli Monitoring components and base agents. Use these codes when running commands. For information on the product codes for Tivoli Monitoring distributed agents and other Tivoli Management Services based agents, see the product documentation.

Table 155. Component product codes for infrastructure components and base monitoring agents

Component	Product code
Endpoint monitoring agent	tm
IBM i monitoring agent	a4
Linux OS monitoring agent	lz
Summarization and Pruning Agent	sy
Tivoli Enterprise Monitoring Server	ms
Tivoli Enterprise Portal browser client	cw
Tivoli Enterprise Portal desktop client	cj
Tivoli Enterprise Portal Server	cq
IBM Tivoli Universal Agent	um
Tivoli Enterprise Monitoring Automation Server	as
IBM Message Service Client Library	w0
UNIX Log Alert monitoring agent	ul
UNIX OS monitoring agent	ux
Warehouse Proxy	hd
Windows OS monitoring agent	nt
Eclipse Help Server	kf
Agentless monitoring for Windows operating systems	r2
Agentless monitoring for AIX operating systems	r3
Agentless monitoring for Linux operating systems	r4
Agentless monitoring for HP-UX operating systems	r5
Agentless monitoring for Solaris operating systems	r6
Business System Manager common agent	r9
Tivoli Performance Analyzer	pa
Tivoli Performance Analyzer Domain for DB2	p0
Tivoli Performance Analyzer Domain for OS agent	p3
Tivoli Performance Analyzer Domain for Oracle	p4
Tivoli Performance Analyzer Domain for System P	p6
Tivoli Performance Analyzer Domain for ITCAM RT	pi

Table 155. Component product codes for infrastructure components and base monitoring agents (continued)

Component	Product code
Tivoli Performance Analyzer Domain for VMware	pu
IBM Tivoli Composite Application Manager Agent for DB2	ud
IBM Tivoli Composite Application Manager Extended Agent for Oracle	rz
Tivoli Enterprise Services User Interface Extensions (tacmd CLI)	ue
IBM GSKit Security Interface	gs
Tivoli Monitoring Shared Libraries	ax
IBM HTTP Server	iu
IBM Tivoli Enterprise Portal Server Extensions	iw
Tivoli Enterprise supplied JRE	jr
Tivoli Enterprise Management Agent Framework Update	tf

A complete, alphabetical list of product codes for IBM Tivoli Monitoring can be found at this Web site: <http://www-304.ibm.com/support/docview.wss?uid=swg21265222&myns=swgtiv&mynp=OCSSZ8F3&mync=E>.

Platform codes required for remote agent deployment

Table 156 lists the platform codes required by the commands related to remote agent deployment.

Table 156. Platform codes required for remote agent deployment

Platform	Code
AIX R4.1	aix4
AIX R4.2	aix42
AIX R4.2.0	aix420
AIX R4.2.1	aix421
AIX R4.3	aix43
AIX R4.3.3	aix433
AIX R5.1 (32 bit)	aix513
AIX R5.1 (64 bit)	aix516
AIX R5.2 (32 bit)	aix523
AIX R5.2 (64 bit)	aix526
AIX R5.3 (32 bit)	aix533
AIX R5.3 (64 bit)	aix536
HP-UX R10.01/10.10	hp10
HP-UX R10.20	hp102
HP-UX R11 (32 bit)	hp11
HP-UX R11 (64 bit)	hp116
HP-UX R11 Integrity (32 bit)	hpi113
HP-UX R11 Integrity (64 bit)	hpi116
Linux Intel R2.2	li622
Linux Intel R2.2 (32 bit)	li6223
Linux Intel R2.4	li624

Table 156. Platform codes required for remote agent deployment (continued)

Platform	Code
Linux Intel R2.4 GCC 2.9.5 (32 bit)	li6242
Linux Intel R2.4 (32 bit)	li6243
Linux Intel R2.4 GCC 2.9.5 (64 bit)	li6245
Linux Intel R2.6 GCC 2.9.5 (32 bit)	li6262
Linux Intel R2.6 (32 bit)	li6263
Linux Intel R2.6 GCC 2.9.5 (64 bit)	li6265
Linux S390 R2.2 (31 bit)	ls322
Linux S390 R2.2 (31 bit)	ls3223
Linux S390 R2.2 (64 bit)	ls3226
Linux S390 R2.4 (31 bit)	ls324
Linux S390 R2.4 GCC 2.9.5 (31 bit)	ls3242
Linux S390 R2.4 (31 bit)	ls3243
Linux S390 R2.4 GCC 2.9.5 (64 bit)	ls3245
Linux S390 R2.4 (64 bit)	ls3246
Linux S390 R2.6 GCC 2.9.5 (31 bit)	ls3262
Linux S390 R2.6 (31 bit)	ls3263
Linux S390 R2.6 GCC 2.9.5 (64 bit)	ls3265
Linux S390 R2.6 (64 bit)	ls3266
Linux AMD64 R2.4 (32 bit)	lx8243
Linux x86_64 R2.4 (64 bit)	lx8246
Linux AMD64 R2.6 (32 bit)	lx8263
Linux x86_64 R2.6 (64 bit)	lx8266
Linux ia64 R2.4 (64 bit)	lia246
Linux ia64 R2.6 (64 bit)	lia266
Linux ppc R2.4 (64 bit)	lpp246
Linux ppc R2.6 (32 bit)	lpp263
Linux ppc R2.6 (64 bit)	lpp266
MVS™	mvs
Digital UNIX	osf1
OS/2	os2
OS/400®	os400
Solaris R2.4	sol24
Solaris R2.5	sol25
Solaris R2.6	sol26
Solaris R7 (32 bit)	sol273
Solaris R7 (64 bit)	sol276
Solaris R8 (32 bit)	sol283
Solaris R8 (64 bit)	sol286
Solaris R9 (32 bit)	sol293
Solaris R9 (64 bit)	sol296

Table 156. Platform codes required for remote agent deployment (continued)

Platform	Code
Solaris R10 (32 bit)	sol503
Solaris R10 (64 bit)	sol506
Solaris R10 Opteron (32 bit)	sol603
Solaris R10 Opteron (64 bit)	sol606
Tandem Itanium (64 Bit)	ta6046
Tandem MIPS (64 Bit)	tv6256
Tru64 V5.0	tsf50
Tivoli Enterprise Monitoring Server support	tms
Tivoli Enterprise Portal Server support	tps
Tivoli Enterprise Portal Desktop client support	tpd
Tivoli Enterprise Portal Browser client support	tpw
Common code, sample libraries and documentation	unix
Windows Itanium 64 bit	WIA64
Windows x86-64 64 bit	WIX64
Windows (all other environments)	winnt

Application support codes

Table 157 lists the various IBM Tivoli Monitoring components and the codes that represent them; these show up when invoking the `cinfo -t` command.

Table 157. Application support codes

Component	Code
Tivoli Enterprise Monitoring Server	TEMS
Tivoli Enterprise Portal Server	TEPS
Tivoli Enterprise Portal	TEP
Windows	WI
Tivoli Enterprise Monitoring Server application support (Windows systems)	WICMS
Tivoli Enterprise Portal Server application support (Windows systems)	WICNS
Tivoli Enterprise Portal browser client application support (Windows systems)	WIXEB
Tivoli Enterprise Portal desktop client application support (Windows systems)	WIXEW
Tivoli Enterprise Monitoring Server application support (Linux or UNIX systems)	tms
Tivoli Enterprise Portal Server application support (Linux or UNIX systems)	tps
Tivoli Enterprise Portal browser client application support (Linux or UNIX systems)	tpw
Tivoli Enterprise Portal desktop client application support (Linux or UNIX systems)	tpd

Appendix E. Agent configuration and environment variables

This appendix contains information about the IBM Tivoli Monitoring configuration files that are preserved during an upgrade, behavior of customized configuration settings, and information about how to create persistent configuration changes.

This appendix also contains a list of environment variables related to the components and monitoring agents of IBM Tivoli Monitoring.

The following table lists the information topics related to IBM Tivoli Monitoring agent configuration and environment variables:

Table 158. Information topics related to IBM Tivoli Monitoring agent configuration and environment variables

Topic	Where to find information
Configuration files that are preserved during an upgrade.	To understand which configuration files are preserved during an upgrade, and which configuration files are not preserved, see “Configuration files preserved during an upgrade.”
Product behavior with customized configuration settings.	For an overview of product behavior regarding customized configuration settings, see “Product behavior with custom configuration settings” on page 890.
Persistent configuration changes.	For information about how to create persistent configuration changes for configuration environment variables, see “Persistent configuration changes” on page 891.
IBM Tivoli Monitoring environment variables.	For a list of environment variables related to the components and monitoring agents of IBM Tivoli Monitoring, see “Environment variables” on page 893.

Configuration files preserved during an upgrade

Custom settings and variables located in configuration files might be lost when upgrading IBM Tivoli Monitoring.

These files include K??CMA.ini and K??ENV files for Windows systems, and *.ini, *.config, k*env files for Linux or UNIX, and more. This section explains which configuration files are preserved during an upgrade, which configuration files are not preserved, and the behavior of custom settings in configuration files.

Preserving user custom settings

This section explains how user custom settings are preserved.

During an upgrade many product files are replaced with newer versions. Other files are merged with existing files to produce the updated version. Still other files are generated by the installation process using values you provide. For more information see, “Product behavior with custom configuration settings” on page 890.

The following general rules apply to how user custom settings are preserved:

- User-defined constructs are kept. For example, situations, policies, queries, and workspaces are always preserved automatically on upgrade.
- Values you can change through a supported product interface are preserved.

- Values that you changed manually (for example, because of a technote or as directed by IBM software support) are probably preserved. Any value that you changed manually that was restored to a default value during an upgrade is recoverable from the backups made during the upgrade process.

The configuration process works with two basic types of files:

- **K??CMA.ini, k*env** (initialization) files are used to collect the inputs from the installation process. This input information is the set of responses to installation questions that are captured as keyword-value pairs. This information is laid down with default values and basic information about the installed components.
- **K??ENV, *.config** (configuration) files are generated from the values in the initialization files and the values entered during configuration.

The initialization files are source files and the configuration files are output files. Although the source files are modified by the configuration tools and sometimes by hand, the configuration output files are rarely modified. Configuration files are generated files, therefore anything changed manually in configuration files is lost during the reconfiguration of that component.

To recover a manual customization after upgrading, perform the following steps:

1.
 - **Windows** Compare the new version of the K??CMA.ini or K??ENV file with the version saved in the itm_home\backup\backups\date_and_time_of_upgrade directory.
 - **Linux** **UNIX** Compare the new version of the *.ini file with the version saved in the same directory as *.ini.bak.
2. Change the installer-supplied defaults to the hand-edited values found in the backup file. Make any changes required to carry your hand-edited custom settings forward, and save the new file.

Files that are preserved on upgrade

The files in this section are preserved when you upgrade your IBM Tivoli Monitoring environment.

If you modified these files or settings, you can expect the changes to be preserved on upgrade:

-D flags in Applet.html

Any changes made to the -D flags are preserved.

Bannerimage.html

Bannerimage.html in the CNB directory. If you add a customer image for your own banner, it is preserved.

-D flags in "cnp" batch or script files

Any changes made to the -D flags on the Java calls of these files are preserved. The specific file names are as follows:

- **Linux** **UNIX** cnp.sh and cnp_inst.sh
- **Windows** cnp.bat and cnp_inst.bat

where *inst* is the name of an instance of the Tivoli Enterprise Portal to connect to. The -D flags in cnp.sh are not preserved.

K??ENV, *.ini

The current settings from ENV files are preserved by checking “key = value” and adding keys that did not exist in the new file from the old file, and replacing the value from the old file in the new file. Keys with default values are preserved.

K??CMA.ini, k*env

If a change was made using a provided configuration tool, the value is always preserved. If you were instructed by a technote or by an IBM support engineer to make a manual change the value is most likely, but not always, preserved (depending on what was changed and why).

OM_TEC.config

In one of the following directories:

- **Linux** **UNIX** itm_home/tables/temsname/TECLIB
- **Windows** itm_home\cms\TECLIB

tecserver.txt

In one of the following directories:

- **Linux** **UNIX** itm_home/tables/temsname/TECLIB
- **Windows** itm_home\cms\TECLIB

jrelevel.js

In one of the following directories:

- **Linux** **UNIX** itm_home/ARCH/cw/. Where ARCH is a specific architecture, such as **li6263**.
- **Windows** itm_home\CNB\

Files that are not preserved on upgrade

This section contains the files that are not preserved when you upgrade your IBM Tivoli Monitoring environment.

If you modified these files or settings, you must make a backup before performing the upgrade. After the upgrade, you must reconstitute your modifications from your backups.

For the Manage Scripts feature in Tivoli Enterprise Portal

The Manage Scripts feature in the Tivoli Enterprise Portal has a set of built-in scripts that manage 3270 terminal session navigation. These built-in functions persist after you install IBM Tivoli Monitoring. However, you must save the custom scripts that you define for the navigation of 3270 terminal sessions. For more information, see the *IBM Tivoli Monitoring Tivoli Enterprise Portal User's Guide*.

For components running on Windows systems OR Linux or UNIX systems

The following IBM Tivoli Enterprise Console event synchronization files are not preserved on upgrade:

- In the TEC_CLASSES directory of the rulebase created during IBM Tivoli Enterprise Console event synchronization installation: omegamon.baroc, Sentry.baroc.
- In the TEC_RULES directory of the rulebase created during IBM Tivoli Enterprise Console event synchronization installation: omegamon.rls.

Before these files are replaced, backup copies are made automatically and placed in the same directories as the original files and have the **.bac** suffix added to their names. You can open these backup files and manually migrate modifications.

Note: The **.baroc** and **.rls** files are backed up only if you choose to automatically upgrade the specified rulebase.

For components running on Windows systems

The following files are not preserved on upgrade:

- On the Tivoli Enterprise Portal Server, the **CNP.bat** and **applet.html** files are built based on the content of the CNPS, CNP, and CNB directories. Except for the -D flags, everything in these files is regenerated.
- The **buildpresentation.bat** file is generated during installation. Any updates to this file are lost.

For components running on Linux or UNIX systems

The following files are not preserved on upgrade:

- **ARCH/cq|cj|cw/***, where ARCH is a specific architecture, such as **li6263** and one of the following two-letter codes: cq, cj, or cw. These files are overwritten during upgrade. One exception to this rule is the **ARCH/cw/jrelevel.js** file, the content of which is preserved.

- The autostart scripts. These files are regenerated during the installation and configuration and have different names based on the platform:
 - Linux, Solaris: `/etc/init.d/ITMAgents*`
 - HP-UX: `/sbin/init.d/ITMAgents*`
 - AIX: `/etc/rc.itm*`

Special exceptions

On Linux or UNIX systems, several IBM Tivoli Monitoring configuration files exist with values that do not typically persist during an upgrade. However, the values in the following files do persist:

- `itm_home/tables/temsname/*.txt`, where ***.txt** refers to all file names that have the text (***.txt**) tag, including **partition.txt** (previously including **glb_site.txt**).

Note: The **glb_site.txt** file is not supposed to be preserved. It is a configuration output file and its values are gathered and stored during normal configuration.

- The following files are generated from **kbbenv.ini** and **ms.ini**:
 - `itm_home/tables/temsname/KBBENV`
 - `itm_home/config/hostname_ms_temsname.config`

Changes made to these files are saved in the same way that changes made to the source .ini files are saved.

Product behavior with custom configuration settings

This section provides an overview of product behavior regarding custom configuration settings.

Before you upgrade, make sure you back up your IBM Tivoli Monitoring environment, especially if you have a more complicated environment. Most companies have a standard process for preserving a backup image of all computer systems, including the IBM Tivoli Monitoring environment.

General operations

Several types of variables control the operation of an IBM Tivoli Monitoring component.

See the following basic types of variables:

User-modified variables

You generate these settings when doing a configuration, changing a value in the CLI or GUI configuration GUIs, and then saving the changes. These variables are user-modifiable variables. They are typically stored in initialization `K??CMA.ini`, `k*env` files on the disk. The output (`K??ENV, *.config`) files use these variables, such as **key=\$VAR1\$**, which are then substituted with the value that you specify.

Static variables

The second type of setting is a variable for an internal component that you cannot configure through the common configuration tools. These variables are stored in `K??CMA.ini`, `K??ENV`, `k*env`, `*.ini` files, in entries such as **NUM_TIMES_TO_TRY=4**. No dynamic substitution takes place. This type of setting is static.

The following list describes persistence in several scenarios:

- During an upgrade from one version of a component to another, the user-modified variables are always persisted with the installation values that are set through the configuration tools.
- If you manually modify these values, for example, adding a new value to the end of an existing **\$VR\$** value, that new value is not persisted.
-

The static values for the system are typically not persisted because they are internal component variables, and not typically exposed to the user for configuration. These static variables are not

documented. In most cases, do not change them. Changing these values without instructions from a technote or IBM Software Support can lead to unpredictable results. To ensure stability of the system, the default values are restored when you upgrade to a new component version level.

Although these variables normally are not persisted, a variable that you add that does not already have a key is persisted. This type of value is normally created under the guidance of IBM Software Support.

Keep notes about any changed static variables that you modify in response to instructions from technotes, IBM Software Support, and so on. After performing an upgrade, examine if the values need to be reapplied.

Persistent configuration changes

Advanced users can apply override values to component customization.

Applying override values ensures that values are retained during an upgrade. You should test this in your environment first before you apply it globally.

Windows persistent configuration changes

The Windows agent process preserves configuration changes by design. Updated variables in the `kxxcma.ini` file, where `xx` denotes the product code, are kept in the Override Local Settings section.

These variables are used during each configuration to update the Windows registry entries that the agents use at run time.

Example:

1. Select the monitoring agent that you want to update in the Manage Tivoli Enterprise Monitoring Services window.
2. Right-click the monitoring agent and select **Advanced > Edit Variables**.
3. Click **Add**.
4. Select **KDC_FAMILIES** from the drop-down list of variables. The `@Protocol@` value is displayed in the **Value** field.
5. Append **ephemeral:y** to the current value. The resulting string is `@Protocol@ ephemeral:y`.
6. Click **OK** to save the change.
7. Click **OK** again.

Linux or UNIX persistent configuration changes

The files that contain configuration information for IBM Tivoli Monitoring Linux or UNIX processes are located in the `itm_home/config` directory.

One such file is the `xx.ini` file where `xx` denotes the product code. For example, `ux` is the product code for the UNIX OS agent. Additional information is stored in locations such as `itm_home/config/.ConfigData/kxxenv` where the user defaults and configuration dialog variables are stored. Some information is also stored in other files such as `env.config`. All these files are processed to create a `.config` file.

Each time an agent is started (or stopped) the agent `.config` files are created in one of the following formats:

- `xx.config`
- `xx_instance.config`
- `hostname_xx_instance.config`

Tivoli Enterprise Monitoring Server .config files are created during the Tivoli Enterprise Monitoring Server configuration process:

```
./itmcmd config -S -t temsname
```

Tivoli Enterprise Monitoring Server .config files have the format: hostname_ms_temsname.config. The agent uses the .config file to define the runtime configuration environment variables.

Making persistent configuration changes

To make persistent configuration changes, create the file `xx.environment` in the config directory. In this file define variables in the **key=value** format.

Variables defined in the `xx.environment` file have override variables from the .config file during agent startup.

Note:

1. If you add a variable to the `xx.environment` file and during a later configuration, change a parameter that is also in the `xx.environment` file, the `xx.environment` file will override your configuration changes. If you use the more robust way of overriding existing variables as described in Example 2, those configuration changes will be included.
2. Future IBM Tivoli Monitoring release levels might change the meaning and defaults of configuration variables. You must be aware of changes made to configuration variables and make appropriate alterations to achieve your intended result.
3. In the worst case scenario, an override might prevent the agent from starting or running successfully. This scenario requires a manual login and correction. Therefore you must always thoroughly test the setup.
4. Some agents use other installation configuration files. The method outlined in this section works only on Linux or UNIX .ini file configuration.
5. You can find values passed to the component process in the `itm_home/logs/xx.env` file.

Example 1: Overriding the CTIRA_HOSTNAME to change the host name section of the Managed

System Name: Create a file in the config directory called `xx.environment` with the following contents:

```
CTIRA_HOSTNAME=myname
```

The contents of `xx.environment` is included in the environment.

Result: When the component starts, the lines in the `xx.environment` file are included and processed after other lines in the .config file.

Example 2: Altering an existing variable KDC_FAMILIES:

The `KDC_FAMILIES` environment variable is particularly challenging because much of the data is created from the .ConfigData/kxxenv dialog file.

For example, suppose that in the existing .config file the `KDC_FAMILIES` variable looks similar to the following variable:

```
export KDC_FAMILIES='ip.pipe port:1918 ip use:n ip.spipe use:n sna use:n HTTP:1920'
```

But you must have `ephemeral:y` added. In this case, the contents of the `xx.environment` file might be similar to the following variable:

```
KDC_FAMILIES=ip.pipe port:1918 ephemeral:y ip use:n ip.spipe use:n sna use:n HTTP:1920
```

A more robust content line might be written as follows:

```
KDC_FAMILIES=ephemeral:y ${KDC_FAMILIES}
```

Create an `xx.environment` file for each product that you need to modify. That way your overrides persists for upgrades in current releases, and in future releases.

Environment variables

This section contains environment variables related to the components and monitoring agents that comprise the base IBM Tivoli Monitoring product.

Common environment variables

The table in this section lists the environment variables that are common to all components.

Table 159. Common environment variables

Variable	Value type	Purpose
KDC_DEBUG	Y or N	Default is N. Diagnosing RPC communications or connectivity problems between the following components: <ul style="list-style-type: none">• The portal server and the monitoring server.• An agent and the monitoring server.• A remote monitoring server and the hub monitoring server.
KDE_DEBUG	Y or N	Default is N. Diagnosing general communication problems, including TCP and SSL setup and connections, local interface discovery, host name resolution, and so on.
KBS_DEBUG	Y or N	Default is N. Consolidates KDC_DEBUG, KDE_DEBUG, and KDH_DEBUG, into one environment variable to simplify the process of collecting diagnostic information.
KDE_TRANSPORT		Change the default port settings for a monitoring agent. For more information, see “Tivoli Monitoring protocol usage and protocol modifiers” on page 435.
KDC_FAMILIES		Change the default port settings for a monitoring agent. For more information, see “Tivoli Monitoring protocol usage and protocol modifiers” on page 435.
KDEB_INTERFACELIST		Contains interface addresses in the order for which these interface addresses should be discovered and used.
KDH_DEBUG	Y or N	Default is N. Diagnosing connectivity problems with the integrated web server.

Table 159. Common environment variables (continued)

Variable	Value type	Purpose
KBB_RAS1	Trace specification string	<p>Default tracing level: KBB_RAS1=ERROR</p> <p>Diagnosing client request problems: KBB_RAS1=ERROR (UNIT:ctsql IN ER) (UNIT:ctdata IN ER)</p> <p>Diagnosing client or Tivoli Enterprise Monitoring Server interaction problems: KBB_RAS1=ERROR (UNIT:ctsql IN ER) (UNIT:ctdata IN ER) (UNIT:ctcmw IN ER) (UNIT:kv4 IN ER)</p> <p>Diagnosing SQL generation problems: KBB_RAS1=ERROR (UNIT:ctsql IN ER) (UNIT:ctdata IN ER) (UNIT:ctreport ALL)</p> <p>Diagnosing login problems: KBB_RAS1=ERROR (UNIT:ctsql IN ER) (UNIT:ctdata IN ER) (UNIT:ctauth ALL)</p>
KBB_RAS1_LOG		Determines the count limit and the maximum files settings.
KBB_SIG1	To generate tracing, core dumps or stacktraces, on error.	<ul style="list-style-type: none"> • Possible values: -trace, -dumpoff, -asyncoff • Default value: blank

Tivoli Enterprise Portal Server environment variables

The table in this section lists the Tivoli Enterprise Portal Server environment variables. For instructions about changing an environment variable, see *Editing your environment configuration settings* in the *IBM Tivoli Monitoring Administrator's Guide*.

Table 160. Tivoli Enterprise Portal Server environment variables

Variable	Value type	Purpose
DSUSER1	Y or N	Part of a set of 9 variables DSUSER1...DSUSER9 that you can set by using the tacmd configureportalserver command. For more information, see the <i>IBM Tivoli Monitoring Administrator's Guide</i> .

Table 160. Tivoli Enterprise Portal Server environment variables (continued)

Variable	Value type	Purpose
KBB_RAS1	Trace specification string	<p>Default tracing level: KBB_RAS1=ERROR</p> <p>Diagnosing client request problems: KBB_RAS1=ERROR (UNIT:ctsql IN ER) (UNIT:ctdata IN ER)</p> <p>Diagnosing client or Tivoli Enterprise Monitoring Server interaction problems: KBB_RAS1=ERROR (UNIT:ctsql IN ER) (UNIT:ctdata IN ER) (UNIT:ctcmw IN ER) (UNIT:kv4 IN ER)</p> <p>Diagnosing SQL generation problems: KBB_RAS1=ERROR (UNIT:ctsql IN ER) (UNIT:ctdata IN ER) (UNIT:ctreport ALL)</p> <p>Diagnosing login problems: KBB_RAS1=ERROR (UNIT:ctsql IN ER) (UNIT:ctdata IN ER) (UNIT:ctauth ALL)</p>
KDH_SERVICEPOINT	Service point string	Specifies an alternate service point string, which labels an ITM process in the IBM Tivoli Monitoring Service Index web page. If not specified, the service point string defaults to <process-owner>.<system-name>_<product-code>.
KD8_VIEWMODEL_THREAD_POOL_SIZE	positive integer greater than 1	<p>It is not recommended to change the value of this parameter. This value specifies the number of threads available to compute DataProvider model data. A value too high or too small can cause the DataProvider to fail. It should be changed only by the IBM Tivoli Monitoring Administrator if the DataProvider is operating in a high traffic environment, that is a lot of users or requests in a small period of time.</p> <p>This value must be set in kfwenv (Windows) or in cq.ini (UNIX/Linux). Follow these steps to change the value of KD8_VIEWMODEL_THREAD_POOL_SIZE:</p> <ol style="list-style-type: none"> 1. Make a backup copy of the cq.ini or kfwenv file. 2. Change the value of KD8_VIEWMODEL_THREAD_POOL_SIZE 3. Restart the Tivoli Enterprise Portal Server. <p>In general, accept the default value of this parameter.</p>

Table 160. Tivoli Enterprise Portal Server environment variables (continued)

Variable	Value type	Purpose
KD8_VM_IMPORT_ID	user ID	<p>Disable the importing of View Models all together or control which users can import View Models.</p> <p>When a Dashboard user logs in, a View Model import request is sent to the dashboard data provider. View Models in the data provider define the characteristics of what data the Dashboards can display. This occurs once and it happens automatically. Any Dashboard update requiring an updated View Model also happens automatically.</p> <p>This variable is optional and not set by default, which means any Dashboard user that has been authenticated can import View Models. To control which user is allowed to do the import, set this variable to a particular ID such as KD8_VM_IMPORT_ID=user1. To disable all users from doing an import of View Models, set KD8_VM_IMPORT_ID=\$nouser@ or any name that you know does not match a Dashboard user ID.</p> <p>KD8_VM_IMPORT_ID is set in the following file:</p> <ul style="list-style-type: none"> On Windows: <ITM_HOME>\CNPS\kfwenv <p>For example: C:\IBM\ITM\CNPS\kfwenv</p> <ul style="list-style-type: none"> On Linux/UNIX: <ITM_HOME>/config/cq.ini <p>For example: /opt/IBM/ITM/config</p> <p>You must restart the Tivoli Enterprise Portal Server after changing the value of KD8_VM_IMPORT_ID.</p>
KFW_CMW_SITUATION_ADMIN_SUPPRESS	Y or N	To prevent an event from appearing in the situation event console after its situation was stopped.
KFW_CMW_SITUATION_FANOUT_SUPPRESS	Y or N	<p>To enable or disable the broadcast of administration events without an origin node (for example, stopped events) to all managed systems connected to the event's thrunode. Default is Y.</p> <p>When KFW_CMW_SITUATION_ADMIN_SUPPRESS=N, you must define KFW_CMW_SITUATION_FANOUT_SUPPRESS as N to clear the stopped event in the Situation Event Console after a restart.</p>
KFW_CMW_SPECIAL_HUB_ENTERPRISE=N	Y or N	Default is N. When set to Y, associates situations to the Tivoli Enterprise Portal Server.
KFW_DATABUS_INPUT_TRACE_IGNORE_HEARTBEAT	Y or N	Default is N. Reduces trace volume by skipping client heartbeat requests when (UNIT:ctdata IN ER) request tracing is used.
KFW_DATABUS_QUERY_VERBOSE	Y or N	Default is N. Tivoli Enterprise Portal Server client side response time and request life cycle tracing.

Table 160. Tivoli Enterprise Portal Server environment variables (continued)

Variable	Value type	Purpose
KFW_INHIBIT_OPSLOG_FROM_RAS	Message ID string	<p>Specifies a list of Message IDs to turn off information messages in the Tivoli Enterprise Server log. For example, to turn off the following KFW1353I message:</p> <p>KFW1353I Take action 'xxxxxxx' executed on <nodename> by <userid> from <ip address></p> <p>You specify the message ID as follows:</p> <p>KFW_INHIBIT_OPSLOG_FROM_RAS=KFW1353I</p> <p>To turn off multiple information messages specify the message ID's as follows:</p> <p>KFW_INHIBIT_OPSLOG_FROM_RAS=KFW1353I KFW1010I KFW1100I</p>
KFW_IOR_IHS_HTTP_PORT	port number	<p>You must set this variable to indicate the redirected port number you want to use. Add the following statement to the cq.ini file associated with the Tivoli Enterprise Portal Server, and then recycle the Tivoli Enterprise Portal Server before re-testing:</p> <p>KFW_IOR_IHS_HTTP_PORT=<port#></p> <p>Replace <port#> with the redirected port number you want to use, for example 8080.</p> <p>Note: This is only one of the variables you need to set in order to reconfigure the HTTP port. For more information, see “Controlling port number assignments” on page 421.</p>
KFW_MIGRATE_CMS	Y or N	Default is N. When set to Y, causes initial migration of managed objects and user IDs from the monitoring server the first time the Tivoli Enterprise Portal Server starts after installation.
KFW_MIGRATE_FORCE	Y or N	Default is N. When set to Y, SQL seed files are processed even if the date, time, or size of the files have not changed since prior seeding. This occurs when running the buildpresentation.bat file (Windows) and when running the InstallPresentation.sh file or reconfiguring the Tivoli Enterprise Portal Server (UNIX/Linux). To be effective, this value must be set in kfwalone (Windows) or in cq.ini or lnxenvnocms (UNIX/Linux).
KFW_MIGRATE_VERBOSE	Y or N	Default is N. When set to Y, provides greater detail about operations performed in migrate.log when seeding the Tivoli Enterprise Portal Server database. This occurs when running the buildpresentation.bat file (Windows) and when running the InstallPresentation.sh file or reconfiguring the Tivoli Enterprise Portal Server (UNIX/Linux). To be effective, this value must be set in kfwalone (Windows) or in cq.ini or lnxenvnocms (UNIX/Linux).

Table 160. Tivoli Enterprise Portal Server environment variables (continued)

Variable	Value type	Purpose
KFW_ORBPARAM	Message ID string	To restrict the portal server to using AES 128 bit encryption algorithms, set the following property in cq.ini: KFW_ORBPARAM=-Dvbroker.security.cipherList=TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA,TLS_ECDH_RSA_WITH_AES_128_CBC_SHA,TLS_RSA_WITH_AES_128_CBC_SHA If the Tivoli Enterprise Portal Server is running on Windows, add this environment variable to the kfwenv file and restart the Tivoli Enterprise Portal Server.
KFW_REPORT_FIND_WAREHOUSE_AT_STARTUP	Y or N	Default is Y. When set to Y, the portal server queries the data warehouse during startup. If set to N, the warehouse will be initialized at the first deliberate query instead of pre-initialized.
KFW_REPORT_ROW_LIMIT	positive integer or 0	Default is 100000. Set this variable to limit the number of rows returned for a report request. A KFWITM574W warning message displays in the view indicating the result set was truncated. Set to 0 to disable. This variable affects only report requests when the user selects the "return all rows" option on the View Properties window.
KFW_REPORT_TERM_BREAK_POINT	Y or N	Default is N. When set to Y, specifies the point where a historical request selects from short-term or long-term history data.
KFW_REPORT_WITH_UR	Y or N	Default is N. When set to Y, it appends "with UR" to the SQL statement against the DB2 Tivoli Data Warehouse or Summarization and Pruning table issued by the Tivoli Enterprise Portal Server. It does not lock out the row and tables in the database; and allows other commands access to the Warehouse database at the same time.
KFW_SQL_VERBOSE	Y or N	Default is N. When set to Y, provides RAS1 trace of each SQL statement issued by the Tivoli Enterprise Portal Server to the Tivoli Enterprise Monitoring Server or to the DB2/ODBC data source. The use of this variable is deprecated because the same capability is available using the standard RAS1 trace setting "KBB_RAS1=ERROR (UNIT:ctsql IN ER)"
MSG_MODE	kms or MSG2	Specify kms if you want to use IBM Tivoli Monitoring Operations Logging. Specify MSG2 to use MSG2 logging.
TEPS_SDA	Y or N	N disables the self-describing agent capability at the portal server, whereas Y enables it. The self-describing agent capability is enabled by default at the portal server. Note: Do not specify YES or NO; instead, always specify Y or N.

Table 160. Tivoli Enterprise Portal Server environment variables (continued)

Variable	Value type	Purpose
TEPS_MANIFEST_PATH	Any valid path	<p>The location where you want the portal server to store the manifest and JAR files it collects from the self-describing agents. This parameter is normally set during component installation. Use the following steps:</p> <ul style="list-style-type: none"> • On Windows: <ol style="list-style-type: none"> 1. On the computer where the monitoring server is installed, click Start → Programs → IBM Tivoli Monitoring → Manage Tivoli Monitoring Services. 2. Right-click Tivoli Enterprise Portal Server and click Advanced → Edit Variables. 3. When you are asked if you want to stop the service, answer Yes. 4. The default file location that is set for TEPS_MANIFEST_PATH is <i>ITMinstall_dir</i>\cnps\support where <i>ITMinstall_dir</i> is the directory where you installed the product. You can change this location if you want to do so. • On Linux or UNIX: <ol style="list-style-type: none"> 1. On the computer where the monitoring server is installed, change to the <install_dir>/config/ directory. 2. Open the cq.config file. 3. The default file location that is set for TEPS_MANIFEST_PATH is <i>ITMinstall_dir</i>/arch/cq/support where <i>ITMinstall_dir</i> is the directory where you installed the product. You can change this location if you want to do so.

Tivoli Enterprise Monitoring Server environment variables

The table in this section lists the Tivoli Enterprise Monitoring Server environment variables. For instructions about changing an environment variable, see *Editing your environment configuration settings* in the *IBM Tivoli Monitoring Administrator's Guide*.

Table 161. Tivoli Enterprise Monitoring Server environment variables

Variable	Purpose
ATTRLIB	Specifies the Tivoli Enterprise Monitoring Server attribute (ATR) files directory.
CANDLE_HOME	The directory where the product was installed.

Table 161. Tivoli Enterprise Monitoring Server environment variables (continued)

Variable	Purpose
CLEARDEPLOYSTATUSFREQ	<p>Specifies a schedule for the periodic removal of completed transactions from the Tivoli Enterprise Monitoring Server's Deployment Status table. Enabling this feature allows you to review completed deployment transactions at opportune times as well as reduce the amount of Tivoli Enterprise Monitoring Server overhead by maintaining a modest table size. You can enable this feature at Tivoli Enterprise Monitoring Server startup by adding the variable to the Tivoli Enterprise Monitoring Server's environment file:</p> <p>CLEARDEPLOYSTATUSFREQ=X</p> <p>Where X = number of hours between the automated clearing of the Deployment Status table. If X is zero (0) or this environment variable is not specified in the configuration file, automatic clearing of the table is disabled. Valid values are zero (0) to 720.</p>
CMS_BINPATH	Specifies the Tivoli Enterprise Monitoring Server binary files directory.
CMS_EXTERNALBROKERS	Specifies whether there are external brokers.
CMS_FTO	Specifies whether to use Tivoli Enterprise Monitoring Server Hot Standby.
CMS_MSGBASE	Applies to is/5 platform Tivoli Enterprise Monitoring Agent only. Specify the MSG2 message file for Tivoli Enterprise Monitoring Agent framework messages.
CMS_NODEID	The Tivoli Enterprise Monitoring Server IBM Tivoli Monitoring Node ID.
CMS_NODE_VALIDATION	Enables the rejection of incorrect managed system names.
DEPLOYQUEUE SIZE	Specifies the maximum number of requests that the request queue on the agent deployment controller.
DEPLOYTHREADPOOLSIZE	Specifies the number of threads that are available to the deployment controller.
DEPOTHOME	Specifies the root directory of the agent depot on the file system.
KDC_GLBSITES	Specifies the Tivoli Enterprise Monitoring Server global sites network file with Tivoli Enterprise Monitoring Server Hub host names.
KDH_SERVICEPOINT	Specifies an alternate service point string, which labels an ITM process in the IBM Tivoli Monitoring Service Index web page. If not specified, the service point string defaults to <process-owner>.<system-name>_<product-code>.
KDS_CATLGLIB	Specifies the catalog library.
KDS_HUB	Specifies that this Tivoli Enterprise Monitoring Server is a Hub Tivoli Enterprise Monitoring Server (*LOCAL) or Remote Tivoli Enterprise Monitoring Server (*REMOTE).
KDS_NCS	Whether to use IBM Tivoli Monitoring network lookup services.
KDS_RULELIB	Specifies the rule library.
KDS_RUN	Specifies Tivoli Enterprise Monitoring Server components or probes to run at Tivoli Enterprise Monitoring Server startup.
KDS_START	Specifies Tivoli Enterprise Monitoring Server KDS component startup.
KDS_VALIDATE	Whether to use Tivoli Enterprise Monitoring Server authentication.

Table 161. Tivoli Enterprise Monitoring Server environment variables (continued)

Variable	Purpose
KGE_EXPIRE_INTERVAL_ENTITY KGE_EXPIRE_INTERVAL_PERMIT	<p>If you set KMS_SECURITY_COMPATIBILITY_MODE=N for strict command security, it means that the sending and receiving system clocks must be synchronized with a delta of less than 75 minutes. Note that this 75 minute restriction does accommodate timezone differential. The restriction only applies to system clocks whose Coordinated Universal Time (UTC) values are off by more than 75 minutes; it does not apply to systems whose local times are different. If the time delta exceeds 75 minutes, a command request will be marked as expired when it reaches the destination, and it will fail to execute.</p> <p>If you have remote agent endpoints whose system clocks are widely divergent from their connecting monitoring server, and you cannot readily change the agent clocks, then you can add these two environment variables to the Tivoli Enterprise Monitoring Server:</p> <ul style="list-style-type: none"> • KGE_EXPIRE_INTERVAL_ENTITY=<i>nnn</i> • KGE_EXPIRE_INTERVAL_PERMIT=<i>nnn</i> <p>Where <i>nnn</i> is the time delta expressed in minutes of system clock differential between your Tivoli Enterprise Monitoring Server and its connecting agent. If, for example, the agent's system clock was 6 hours ahead of the Tivoli Enterprise Monitoring Server system clock, you would add these two environment variables to your Tivoli Enterprise Monitoring Server and set the minute values to a number greater than 360 (60 * 6). The minute values in both environment variables should be the same.</p>
KGLCB_FSYNC_ENABLED	Whether to enable UNIX or Linux fsync calls.
KGL_CBTBUFCNT	Specifies Tivoli Enterprise Monitoring Server internal table buffer count.
KGL_CBTBUFSZ	Specifies Tivoli Enterprise Monitoring Server internal table buffer size.
KGL_KEYRING_FILE	Specifies LDAP authentication SSL GSKit keyring file.
KGL_KEYRING_LABEL	Specifies LDAP authentication SSL GSKit keyring label.
KGL_KEYRING_PASSWORD	Specifies LDAP authentication SSL GSKit keyring password.
KGL_KEYRING_STASH	Specifies LDAP authentication SSL GSKit keyring password stash file.
KGL_KGLMSGBASE	Specifies the Tivoli Enterprise Monitoring Server KGL message file directory.
KGL_LDAP_BASE	LDAP authentication LDAP search base.
KGL_LDAP_BIND_ID	LDAP authentication LDAP server bind Distinguished Name (DN).
KGL_LDAP_BIND_PASSWORD	LDAP authentication LDAP server bind password.
KGL_LDAP_HOST_NAME	LDAP authentication LDAP server host name.
KGL_LDAP_PORT	LDAP authentication LDAP server port.
KGL_LDAP_SSL_ENABLED	Whether to use LDAP SSL communications.
KGL_LDAP_USER_FILTER	LDAP authentication user filter.
KGL_LDAP_VALIDATE	Whether to use LDAP authentication.
KGL_MSG2_EVENTLOG	Whether to format Event Log.
KGL_MSG2_UNIVERSAL	Whether to enable Universal Messages.
KGL_TRC1	Whether to enable the error log.
KHD_HISTRETENTION	Specifies the default retention period in hours for the short-term history files (default is 24 hours). This can be used to reduce the amount of data kept on disk after a successful upload to the warehouse is performed.

Table 161. Tivoli Enterprise Monitoring Server environment variables (continued)

Variable	Purpose
KHD_HISTRETENTION	Specifies the default retention period in hours for the short-term history files (default is 24 hours). This can be used to reduce the amount of data kept on disk after a successful upload to the warehouse is performed.
KIB_MAXCOLS	Tivoli Enterprise Monitoring Server internal dictionary column maximum.
KMS_DISABLE_TEC_EMITTER	TEC Emitter to be disabled.
KMS_EVAL_REFLEX_AT_TEMS	When set to Y, this variable maximizes your level of duper optimization. You might want to increase your level of duper optimization if you have a large amount of reflex actions executing at your managed systems (agents). At each monitoring server in your environment, set KMS_EVAL_REFLEX_AT_TEMS=Y to cause the evaluations of the actions to be handled by duper on the monitoring server, while the actions are still directed to the managed systems.
KMS_OMTEC_GLOBALIZATION_LOC	TEC Integration Globalization locale.
KMS_OMTEC_INTEGRATION	TEC Integration enabled.

Table 161. Tivoli Enterprise Monitoring Server environment variables (continued)

Variable	Purpose
KMS_SECURITY_COMPATIBILITY_MODE	<p>Monitoring agents running the Tivoli Enterprise Management Agent framework at version 6.3 or later use security tokens to validate that the request originated from a V6.3 or later monitoring server or portal server in your Tivoli Monitoring environment. The monitoring agents also extract the identity of the user who initiated the request from the security token so that the User ID can be included in an audit record. The 32 character encryption key that is configured for each monitoring component is used to encrypt and decrypt the security token and must be the same for all components in your monitoring environment.</p> <p>The KMS_SECURITY_COMPATIBILITY_MODE variable (Default=Y) is intended to help with customer migration to Tivoli Monitoring V6.3 by allowing a pre-V6.3 component to initiate a take action and execute command request that does not include a security token. If the Y value is set, the Tivoli Enterprise Monitoring Server will generate a default security token so that the take action and execute command request can still be processed successfully at the monitoring agent endpoints connected to that monitoring server.</p> <p>KMS_SECURITY_COMPATIBILITY_MODE=Y also ensures that Take Action and execute command requests can be processed if the system clocks are not synchronized to Coordinated Universal Time (UTC) across the systems with monitoring servers, the portal server, and monitoring agents.</p> <p>KMS_SECURITY_COMPATIBILITY_MODE=Y also has another property. If local cryptographic functions are not available, for example, because the Integrated Cryptographic Services Facility (ICSF) was not configured on a z/OS monitoring server, then that monitoring server would in effect disable security token checking and still allow commands to be executed at the monitoring agent endpoints. This property of the 'Y' setting is intended to give you time to implement cryptographic support on all of your IBM Tivoli Monitoring systems.</p> <p>However, if the value is set to N, the Tivoli Enterprise Monitoring Server will not generate a default security token for the monitoring agents connected to it. Therefore, if a monitoring agent is running V6.3 or later of the Tivoli Enterprise Management Agent framework, it will reject a Take Action or execute command request that originated from a monitoring server or portal server that did not generate a security token or if the security token has expired because clocks are not synchronized. The default expiration time for a security token is 75 minutes. If you want to change the expiration time for the security token, change the values of the KGE_EXPIRE_INTERVAL_ENTITY and KGE_EXPIRE_INTERVAL_PERMIT environment variables.</p> <p>It is recommended that once all of your IBM Tivoli Monitoring server components have been migrated to V6.3, the clocks are synchronized to Coordinated Universal Time (UTC) across your monitoring server, portal server, and monitoring agent components, and cryptographic functions are available on all of your monitoring servers, you should set this environment variable to N. This will ensure that every take action and execute command request originates from a trusted source that has provided a valid security token.</p>

Table 161. Tivoli Enterprise Monitoring Server environment variables (continued)

Variable	Purpose
KPX_WAREHOUSE_LOCATION	Allows a fixed warehouse route for the agents connected to that Tivoli Enterprise Monitoring Server when the usage of the Global Location Broker default algorithm is not supported. It is a list of fully qualified, semicolon delimited network names such as: KPX_WAREHOUSE_LOCATION= family_protocol: #network_address[port_number]; ...
KPX_WAREHOUSE_REGCHK	Number of minutes to wait between rechecking the Global Location Broker for any warehouse proxy agent registration changes. The default value is set to 60 minutes.
KSH_DIRECTORY	Specifies the Tivoli Enterprise Monitoring Server SOAP Server HTML files directory.
KT1_TEMS_SECURE	YES is required for the tacmd commands, putfile , getfile , and executecommand . Also required when executeaction is issued with advanced options. Not required when issuing executeaction to agent without advanced options.
MSG_MODE	Specify kms if you want to use IBM Tivoli Monitoring Operations Logging. Specify MSG2 to use MSG2 logging.
NLS1_LOCALEDIR	The directory of the locale file.
RKDSCATL	Specifies the Tivoli Enterprise Monitoring Server catalogue (CAT) files directory.
SQLLIB	Specifies the Tivoli Enterprise Monitoring Server seeding and query (SQL) files directory.
TIMEOUT	Specifies the time in seconds that Agent Deployment tool has to complete a task. If the tool does not complete in the task in the time specified by the TIMEOUT value, the task is terminated. The default value is 600 seconds.
KMS_SDA	N disables the self-describing agent capability and Y enables it. By default, the self-describing agent capability is disabled at the hub monitoring server. Self-describing agent environment variables are enabled by default for remote monitoring servers, portal servers, and self-describing agent enabled agents. However, these components disable self-describing agent capability if connected to a hub monitoring server that has the self-describing agent capability disabled. Enabling or disabling the self-describing agent capability at the hub monitoring server controls the capability across all components, but you can disable the capability individually at a remote monitoring server, portal server, or agent. You must restart the hub monitoring server for a change in the KMS_SDA variable to take effect. Note: Do not specify YES or NO; instead, always specify Y or N.
KMS_SDM_HUB_WAIT_TIMEOUT	The number of seconds that the Remote Tivoli Enterprise Monitoring Server waits for the hub monitoring server to complete a self-describing agent installation request initiated by this Remote Tivoli Enterprise Monitoring Server. If the hub monitoring server does not complete the self-describing agent installation request within this time interval, the Remote Tivoli Enterprise Monitoring Server gives up waiting for completion and returns a status of KFASDM_ST_Server_TimedOut to the agent. One or more self-describing agents will continue to retry the self-describing agent installation for this product until the hub monitoring server product installation is confirmed. A valid value is a number in seconds greater than or equal to 180. The default value is 180 seconds. This parameter is an internal parameter that you might use, in a rare situation, to fine-tune the self-describing agent capability. In general, accept the default value of this variable.

Table 161. Tivoli Enterprise Monitoring Server environment variables (continued)

Variable	Purpose
KMS_TAPPLOG_MAX_QUEUED_EVENTS	<p>Do not change this value unless directed to do so by IBM Software Support. This variable determines the maximum number of application support installation event notifications that are stored in the TAPPLOGT notification queue. Modifying this value might result in a loss of self-describing agent installation change notices and unpredictable self-describing agent results. A valid value is a number greater than or equal to 20. The default value is 100.</p> <p>This parameter is an internal parameter that you might use, in a rare situation, to fine-tune the self-describing agent capability. In general, accept the default value of this variable.</p>
KMS_CONFIG_PROP_USE_CACHE	<p>Do not change this value unless directed to do so by IBM Software Support. When set to Y, this variable directs the Tivoli Enterprise Monitoring Server self-describing agent installation manager to use the local Tivoli Enterprise Monitoring Server internal cache mechanism of the application support installation records table to speed up self-describing agent operations. When set to N, operations against the application support installation records table use only the TEMS KEY1 database I/O routines and bypass the table internal cache. This might slow down self-describing agent operations.</p> <p>This parameter is an internal parameter that you might use, in a rare situation, to fine-tune the self-describing agent capability. In general, accept the default value of this variable.</p>
SOAP_IS_SECURE	<p>Enabling this variable requires all users who submit CT_EMail or CT_Export requests to know the monitoring server credentials. By default this variable is disabled.</p>
TEMS_JAVA_BINPATH	<p>This variable locates the JAVA installation path within the z/OS USS environment. For more information, see <i>Configuring the Tivoli Enterprise Monitoring Server on z/OS</i> (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.omegamon_share.doc_6.3/ztemsconfig/ztemsconfig.htm).</p>

Table 161. Tivoli Enterprise Monitoring Server environment variables (continued)

Variable	Purpose
TEMS_MANIFEST_PATH= <i>file_loc</i>	<p>The location where you want the monitoring server to store the manifest and JAR files it collects from the self-describing agents. You must set this parameter to enable the self-describing capability. This parameter is normally set during component installation.</p> <p>Use the following steps to set the location where you want the monitoring server to store the manifest and JAR files:</p> <ul style="list-style-type: none"> On Windows: <ol style="list-style-type: none"> On the computer where the monitoring server is installed, click Start → Programs → IBM Tivoli Monitoring → Manage Tivoli Monitoring Services. Right-click Tivoli Enterprise Monitoring Server and click Advanced → Edit Variables. When you are asked if you want to stop the service, answer Yes. The default file location that is set for TEMS_MANIFEST_PATH is <i>ITMinstall_dir</i>\CMS\support where <i>ITMinstall_dir</i> is the directory where you installed the product. You can change this location if you want to do so. On Linux or UNIX: <ol style="list-style-type: none"> On the computer where the monitoring server is installed, change to the <i>ITMinstall_dir</i>/config/ directory. Open the <i>tems_hostname_ms_tems_name.config</i> file. The default file location that is set for TEMS_MANIFEST_PATH is <i>ITMinstall_dir</i>/tables/<TEMS_NAME>/support where <i>ITMinstall_dir</i> is the directory where you installed the product. You can change this location if you want to do so.

Tivoli Enterprise Monitoring Automation Server environment variables

The table in this section lists the Tivoli Enterprise Monitoring Automation Server environment variables. For instructions about changing an environment variable, see *Editing your environment configuration settings* in the *IBM Tivoli Monitoring Administrator's Guide*.

Table 162. Tivoli Enterprise Monitoring Automation Server environment variables

Variable	Description	Valid values	Default value
KAS_DEBUG	Enables additional service provider tracing for troubleshooting.	<ul style="list-style-type: none"> N for no troubleshooting messages. Y for some tracing. A for all troubleshooting tracing. D for detailed troubleshooting tracing. M for maximum troubleshooting tracing. 	N

Table 162. Tivoli Enterprise Monitoring Automation Server environment variables (continued)

Variable	Description	Valid values	Default value
KAS_HOSTNAME	The hostname of the Tivoli Enterprise Monitoring Automation Server is automatically discovered and used for URL specification. If the discovered hostname is incorrect, KAS_HOSTNAME can be used to specify the correct fully qualified hostname.		
KAS_REGISTRY_SERVICES_TIMEOUT_GET	Sets the HTTP/S GET timeout.	15 - 900 seconds	120 seconds
KAS_REGISTRY_SERVICES_TIMEOUT_POST	Sets the HTTP/S POST timeout.	15 - 900 seconds	120 seconds
KAS_REGISTRY_SERVICES_TIMEOUT_DELETE	Sets the HTTP/S DELETE timeout.	15 - 900 seconds	120 seconds
KAS_REGISTRY_SERVICES_URL_PROTOCOL	The protocol of the Registry Services. Available protocols must be specified as follows: <ul style="list-style-type: none"> • =HTTP • =HTTPS 	HTTP HTTPS	
KAS_REGISTRY_SERVICES_URL_HOSTNAME	The hostname of the Registry Services.		
KAS_REGISTRY_SERVICES_URL_HTTP_PORT	The HTTP port of the Registry Services.		
KAS_REGISTRY_SERVICES_URL_HTTPS_PORT	The HTTPS port of the Registry Services.		
KAS_SECURITY_SERVICES_ENABLED	Enable Jazz for Service Management Security Services access. You must set KAS_SECURITY_SERVICES_ENABLED to Yes to enable single sign-on support in the Open Services Lifecycle Collaboration Performance Monitoring (OSLC-PM) service provider. When Security Services support is enabled, the service provider connects to the Jazz for Service Management Security Services application to validate LTPA tokens included in HTTP request headers. The service provider does not support basic authentication requests. For the current release of IBM Tivoli Monitoring, the Jazz for Service Management Registry Services and Security Services applications must be installed on the same host, use the same port, username and password, and support HTTPS connections.	Yes, No	No

Tivoli Data Warehouse environment variables

The Warehouse Proxy agent parameters can be changed in the KHDENV file on Windows systems and the hd.ini file in systems other than Windows systems to customize tracing for the Warehouse Proxy agent.

The Summarization and Pruning agent parameters can be changed in the in the KSYENV file on Windows systems and the SY.ini file on non-Windows systems to customize tracing for the Summarization and Pruning Agent.

Table 163. Warehouse proxy environment variables

Variable	Purpose
CTIRA_NCSLISTEN	The number of RPC threads.
KHD_BATCH_USE	Allows inserts to be batched together, reducing the number of transactions against the database, improving performance and reducing network traffic. Default is Y and can be edited by GUI or CLI.
KHD_CNX_POOL_SIZE	The total number of preinitialized ODBC connection objects available to the work queue export threads. The default value should match the KHD_EXPORT_THREADS value.
KHD_CNX_WAIT	The time in minutes to wait before trying to reconnect. Default is 10 minutes.
KHD_CNX_WAIT_ENABLE	A time to wait before a retry. Default is Y. Changing this variable to N does not wait before retries. Setting this variable to N can generate a large log file if the tests to the database fail at each retry. It is advisable not to change this variable.
KHD_EXPORT_THREADS	The number of worker threads exporting data to the database. The default value should match the KHD_CNX_POOL_SIZE value.
KHD_QUEUE_LENGTH	The length of the KHD work queue. This integer identifies the maximum number of export work requests that can be placed on the work queue before the queue starts rejecting requests. The default value of KHD_QUEUE_LENGTH is 1000. Setting this value to 0 means the queue length has no limit.
KHD_SRV_STATUSTIMEOUT	The time in seconds set by default to 600s = 10 minutes. Set KHD_SRV_STATUSTIMEOUT less than KHD_STATUSTIMEOUT by at least 60 seconds. It is advisable not to change this variable.
KHD_WAREHOUSE_TEMS_LIST	A space or coma-delimited separated list of monitoring servers. A warehouse proxy agent serves all the agents reporting to those monitoring servers listed. The same monitoring server name must not appear more than once in all the warehouse proxy Tivoli Enterprise Monitoring Server lists in the entire enterprise environment.
KHD_DB_COMPRESSION	Used to change the configuration of the Warehouse Proxy Agent to enable or disable database compression.
KHD_WHLOG_ENABLE	Used to change the configuration of the Warehouse Proxy Agent to enable or disable data warehouse log tables. This variable is set to N by default.
KHD_SERVER_DIST_COMPRESSION_ENABLE	Used to send distributed agent compressed data. This variable is set to Y by default.
KHD_SERVER_Z_COMPRESSION_ENABLE	Used to allow z/OS clients to send compressed data. This variable is set to N by default.
KHD_CLIENT_COMPRESSION_ENABLE	Specifies that historical data is not compressed if set to N, which is true even if the Warehouse Proxy Agent server has allowed the compression. This variable does not exist by default.
KHD_DEFAULT_TABLE_CONTAINER	Name of the database container where tables should be placed.
KHD_DEFAULT_INDEX_CONTAINER	Name of the database container where indices should be placed.
KHD_PARTITION	Controls whether partitioning is enabled (Y) or disabled (N). When disabled, non partitioned tables and indices are created and the old pruning behavior is in effect. When enabled, partitioned tables and indices are created (if supported by the database) and partitions are maintained.

Table 163. Warehouse proxy environment variables (continued)

Variable	Purpose
KHD_PARTITIONS_UPWARD	Number of partitions to maintain after the current date. This allows for an outage to occur and not cause disruption on the monitoring system, permitting agents to upload their collected historical data. Minimum value of 1, maximum value of 10. Negative values and special characters are not allowed.

Table 164. Summarization and Pruning agent environment variables

Variable	Purpose
KSY_BLACKOUT	Specifies a comma-separated list of exception times where the Summarization and Pruning agent should not start when using the flexible scheduling. The values in the list should be in the format HH:MM-HH:MM where HH must be between 00 and 23 and MM must be between 00 and 59. The starting time must be smaller than the end time of the exception period. For example, to block the Summarization and Pruning agent from starting between 1 and 2 AM and 5 and 6 PM, use the following time statement: 01:00-01:59,17:00-17:59.
KSY_CNP_SERVER_HOST	Tivoli Enterprise Portal Server connection default host.
KSY_CNP_SERVER_PORT	Tivoli Enterprise Portal Server connection default port.
KSY_DAY_AGE_UNITS	The minimum age of data for daily data in days before aggregation is done.
KSY_EVERY_N_DAYS	How often to run the schedule in days (number).
KSY_EVERY_N_MINS	Indicates the frequency between Summarization and Pruning agent runs when using the flexible scheduling. Must be a multiple of 15 minutes, with a minimum of 15 and a maximum of 1440 (one day).
KSY_FIXED_SCHEDULE	Indicates whether the Summarization and Pruning agent is configured for fixed schedule if set to Y or flexible schedule when set to N.
KSY_HOUR_AGE_UNITS	The minimum age of data for hourly data in hours before aggregation is done.
KSY_HOUR_AM_PM	Whether to run in the AM or PM (AM/PM).
KSY_HOUR_TO_RUN	The hour of the day to run. Valid values are 1-12.
KSY_MAX_ROWS_PER_TRANSACTION	Maximum rows per database transaction.
KSY_MAX_WORKER_THREADS	Maximum number of simultaneous worker threads. Default is 1. Recommended value is the number of processors on your server minus 1.
KSY_MINUTE_TO_RUN	The minute of the day to run (number).
KSY_SHIFT1_HOURS	A comma-separated list of off-peak hour numbers for the shift.
KSY_SHIFT2_HOURS	A comma-separated list of peak hour numbers for the shift.
KSY_SHIFTS_ENABLED	Shift periods. Only two shifts are allowed. If shifts are enabled, each hour (0-23) must be specified once.
KSY_START_OF_WEEK_DAY	Start of the week day, for example: 0 = Sunday 1 = Monday.
KSY_TIMEZONE_IND	Time zone indicator. AGENT specifies use the time zone offset of the agent. WAREHOUSE specifies use the time zone offset of the warehouse.
KSY_VACATION_DAYS	Vacation days in a comma-separated list of days in YYYYMMDD format.
KSY_VACATIONS_ENABLED	Whether vacation is enabled with Y or N.

Table 164. Summarization and Pruning agent environment variables (continued)

Variable	Purpose
KSY_WAREHOUSEAGGREGLOG_PRUNE	Use to specify the pruning for the WAREHOUSEAGGREGLOG table. The format of the value should be <i>number.unit</i> , where <i>number</i> is the number of units to retain and <i>unit</i> specifies the retention unit, which can be one of day, month or year.
KSY_WAREHOUSELOG_PRUNE	Use to specify the pruning for the WAREHOUSELOG table. The format of the value should be <i>number.unit</i> , where <i>number</i> is the number of units to retain and <i>unit</i> specifies the retention unit, which can be one of day, month or year.
KSY_WAREHOUSE_DRIVER	The Warehouse Database Connection JDBC Driver.
KSY_WAREHOUSE_PASSWORD	The Warehouse Database Connection Password (encrypted).
KSY_WAREHOUSE_URL	The Warehouse Database Connection JDBC URL.
KSY_WAREHOUSE_USER	The Warehouse Database Connection User (encrypted).
KSY_WEEKENDS_AS_VACATIONS	Whether weekends are vacation with Y or N.
KSZ_JAVA_ARGS	Java arguments.
KSY_WHLOG_ENABLE	Used to change the configuration of the Summarization and Pruning Agent to enable or disable data warehouse log tables. The default value is N.
KSY_DB_COMPRESSION	Used to change the configuration of the Summarization and Pruning Agent to enable (Y) or disable (N) database compression.
KSY_DEFAULT_TABLE_CONTAINER	Name of the database container where tables should be placed.
KSY_DEFAULT_INDEX_CONTAINER	Name of the database container where indices should be placed.
KSY_PARTITION	Controls whether partitioning is enabled (Y) or disabled (N). When disabled, non partitioned tables and indices are created and the old pruning behavior is in effect. When enabled, partitioned tables and indices are created (if supported by the database) and partitions are maintained.
KSY_PARTITIONS_UPWARD	Number of partitions to maintain after the current date. This allows for an outage to occur and not cause disruption on the monitoring system, permitting agents to upload their collected historical data. Minimum value of 1, maximum value of 10. Negative values and special characters are not allowed.
KSY_GROUP_BY_THRESHOLD	Number of past days used in the Summarization and Pruning agent discovery query. This is used to improve the Summarization and Pruning agent runtime and to limit the number of partition locks placed on the table.

Tivoli Monitoring Service Console variables

The table in this section lists the Service Console variables to control account lock-out in the event of multiple failed logon attempts.

Table 165. Service Console environment variables

Variable	Purpose
KDHS_MAX_FAILED_LOGINS	Tivoli Monitoring Service Console lock-out is a process that counts the number of failed logon attempts to the Service Console. This helps to protect the Service Console from attacks that try thousands of different user IDs and passwords. When a Service Console logon fails, the source IP address is recorded along with the user ID and the fail count associated with that IP address is incremented. When the failure count for a specific IP address exceeds KDHS_MAX_FAILED_LOGINS, that IP address is locked out from further login attempts to the Service Console. Default is 5.
KDHS_ACCOUNT_LOCKOUT_PERIOD	When the number of failed logon attempts from an IP address to the Service Console exceeds the number defined in KDHS_MAX_FAILED_LOGINS, that IP address is locked out from further login attempts for the period of time defined in KDHS_ACCOUNT_LOCKOUT_PERIOD. The unit of measure for this variable is seconds. Default is 1800.

Common agent environment variables

The table in this section lists the common agent environment variables that are used by all IBM Tivoli Monitoring agents.

Table 166. Common agent environment variables

Variable	Purpose
CMS_MSGBASE	Applies only to the IBM i platform. Specifies the MSG2 message file for agent framework messages.
CT_CMSDIRECT	<i>Obsolete.</i> Replaced by IBM Tivoli Monitoring V6 firewall communications services. Specify the full NAT firewall address of the monitoring server to connect to, including <i>protocol:address[port#]</i> .
CT_CMSLIST	Required variable that specifies the primary or secondary Tivoli Enterprise Monitoring Server the agent must connect with. Takes a list of monitoring server names in the form <i>network_protocol:hostname</i> or <i>network_protocol:address</i> and delimited by semicolons.
CTIRA_CELL_NAME	<i>Obsolete.</i> Replaced by the CT_CMSLIST agent configuration variable.
CTIRA_HEARTBEAT	The interval, in minutes, of the agent-to-monitoring server heartbeat data sample. The default value is 10 minutes. Shorter heartbeat intervals increase network traffic between the agent and the Tivoli Enterprise Monitoring Server.
CTIRA_HIST_DIR	Required variable that specifies the directory where agent-based short-term history data files are stored. Does not apply to the Tivoli Enterprise Monitoring Server's short-term history data files. The directory must already exist.
CTIRA_HOSTNAME	Used by many, but not all, agents to provide an alternate host name qualifier (<i>subsystem:hostname:nodetype</i>) for the published agent managed system name. Used to remove a long network domain name (such as acme.us.com) from the default agent hostname. Not honored by all agents. For some agents, CTIRA_HOSTNAME might cause unpredictable results in the Tivoli Enterprise Portal navigation tree.
CTIRA_IP_PORT	<i>Applies to z/OS agents only. Do not modify.</i> Set this variable to 0 so the operating system can provide the agent RPC listen port, which prevents a port conflict for some z/OS configurations.

Table 166. Common agent environment variables (continued)

Variable	Purpose
CTIRA_LOG_PATH	Required variable that specifies the directory where the agent's Operations Log file is stored. The file names themselves use the suffixes .LG0 and .LG1. This variable does not apply to agents running on a z/OS system.
CTIRA_MAX_RECONNECT_TRIES	<i>Obsolete.</i> Number of consecutive unsuccessful attempts the agent makes to connect to a Tivoli Enterprise Monitoring Server before giving up and exiting. The default value of 0 means that the agent remains started regardless of its connection status with the monitoring server.
CTIRA_NCSLISTEN	Number of RPC listen server threads to create for the agent. The default value is 10.
CTIRA_NODETYPE	Supplies the agent node type qualifier (<i>subsystem:hostname:nodetype</i>) of the agent managed system name (<i>msn</i>). Provide the agent product indicator in this name. This value is set by the application and should not be set by the user.
CTIRA_OS_INFO	Overrides the default value for agent entries in the Tivoli Enterprise Monitoring Server's "ManagedSystem.Host_Info" column. This variable is used to build the navigation tree for the Tivoli Enterprise Portal Server. The value must match an existing entry in the CNPS/osnames file. This variable is not applicable to subnode type records in the monitoring server's ManagedSystem table.
CTIRA_PRIMARY_FALLBACK_INTERVAL	Forces the agent to switch from the primary Tivoli Enterprise Monitoring Server to one of the defined secondary monitoring servers, because the primary monitoring server is offline or due to network connectivity issues. You want the agent to switch back to the primary monitoring server as soon as possible when it becomes available. This parameter controls the frequency with which the agent performs a lookup of the primary monitoring server. If the primary monitoring server is found, the agent disconnects from the secondary monitoring server and reconnects to the primary monitoring server. A value of zero disables this feature. The minimum value must be 2.5 times CTIRA_RECONNECT_WAIT value. The default value is 4500 seconds, or 75 minutes.
CTIRA_PRODUCT_SEP	Supplies an alternate qualifier for the agent's managed system name (<i>msn</i>). The default value is a colon (:).
CTIRA_RECONNECT_WAIT	Time interval, in seconds, that the agent waits between attempts to register with a Tivoli Enterprise Monitoring Server. The default value is 600 seconds.
CTIRA_REFLEX_ATOMIC	For subnode targets only. Evaluates the situation state by any existing specified display item column name when deciding which reflex situation automation command the agent should execute. Not applicable to reflex situation commands executed or evaluated by the Tivoli Enterprise Monitoring Server. Disable by setting to N. The default value is Y.
CTIRA_REFLEX_TARGET	For subnode targets only. Evaluates situation state by subnode name value in the ORIGINNODE column when deciding which reflex situation automation command the agent should execute. Not applicable to reflex situation commands executed or evaluated by the Tivoli Enterprise Monitoring Server. Disable by setting to N. The default value is Y.
CTIRA_SIT_CLEAN	Number of seconds between garbage cleanup of stale entries in the agent persistent situation file. The default value is 900 seconds.
CTIRA_SIT_FILE	Specifies an alternate name for the default agent-based persistent situation file. This variable should be used only in exceptional conditions because the file name reflects the agent's managed system name. Unsupported for z/OS-based agents.

Table 166. Common agent environment variables (continued)

Variable	Purpose
CTIRA_SIT_MGR	Specifies whether or not to use the agent's persistent situation file when registering with the Tivoli Enterprise Monitoring Server. Using this file improves performance of the monitoring server. Set this variable to N to disable usage. For a z/OS agent, the value must be N, because this feature is not implemented for a z/OS-based monitoring server. For all other platforms, the default value is Y.
CTIRA_SIT_PATH	Required variable that specifies the directory where the agent-based persistent situation file is stored. This agent-only file contains a copy of the Tivoli Enterprise Monitoring Server monitoring situations for the agent's use while registering with the monitoring server. The file is named <code>psit_msn.str</code> , where <i>msn</i> is the managed system name of the agent process. Unsupported for z/OS-based agents.
CTIRA_SUBSYSTEM_ID	Optional variable that overrides the subsystem ID qualifier (<i>subsystem:hostname:nodetype</i>) of the agent managed system name (<i>msn</i>). Describes a monitored resource instance to help make this name unique. Value may also be set by the agent itself.
CTIRA_SYSTEM_NAME	Sets an alternate system name for agent entries in the Tivoli Enterprise Monitoring Server's ManagedSystem.Host_Address column within the <code><NM>mysystem</NM></code> tags. Used to build the Tivoli Enterprise Portal Server's navigation tree. Not applicable to subnode type records in the monitoring server's ManagedSystem table. The maximum allowable length of the value used for this variable is 31 characters.
CTIRA_THRESHOLDS	Specifies the fully qualified name of the XML-based adaptive (dynamic) threshold override file. The default file is located in <code>\$CANDLE_HOME/localconfig/pc/pc_thresholds.xml</code> , where <i>pc</i> is the agent product code. On z/OS systems, the default file name is <i>pc</i> THRES.
IRA_ADAPTIVE_THRESHOLD_MODE	<p>Specifies the adaptive (dynamic) threshold operation mode, either CENTRAL or LOCAL. In CENTRAL mode, threshold overrides are centrally created and distributed to the agent. Do not modify the override XML file, which is the default mode.</p> <p>In LOCAL mode, central distribution to the agent is inhibited, and threshold overrides are locally created and managed. Use LOCAL mode to specify that the agent should ignore enterprise distribution; its affinity will not be registered, so the Tivoli Enterprise Portal cannot override the agent's managed system node. Use this mode cautiously because it causes the Tivoli Enterprise Monitoring Server's thresholds and the agent's thresholds to be out of sync.</p> <p>You must create and manually write override definitions in the same file that is created in CENTRAL mode: <i>managed-system-name_product-code_thresholds.xml</i>. For instance, on Windows, this file is named <code>Primary_myagent_NT_thresholds.xml</code>; on Linux, <code>myagent_LZ_thresholds.xml</code>; on UNIX, <code>myagent_UX_thresholds.xml</code>. On Windows, the file is stored in the <code>%CANDLEHOME%\TMAITM6</code> directory; on Linux and UNIX, the file is stored in <code>\$CANDLEHOME/interp/product-code/bin</code>.</p> <p>The names of the columns to be used when specifying overrides is taken from the attributes file. The override name and objectId must be unique in the XML file. Timestamp is not required.</p> <p>If later you switch back from LOCAL mode to CENTRAL mode, centrally located overrides will again override the local definitions.</p>

Table 166. Common agent environment variables (continued)

Variable	Purpose
IRA_AUTONOMOUS_LIMIT	Sets the saved event limit for autonomous operation. If the specified value is a number (for example, 500), then it is the maximum number of situation event records to be saved by the agent. If the specification is in common disk space units such as KB, MB, or GB (for example, 5 MB), then it is the total amount of disk space to be used by the agent for saving situation event data. The default value is 2 MB.
IRA_AUTONOMOUS_MODE	Turns on (YES) or off (NO) agent autonomous mode. While in autonomous mode, the agent continues to run Enterprise situations. The situation event data persists on disk even after agent restart. On reconnection to the Tivoli Enterprise Monitoring Server, the agent uploads saved situation event data to the monitoring server. The default value is Y.
IRA_DEBUG_AUTONOMOUS	Turns on (Y) or off (N) debug trace for agent autonomous operation. This variable is used for diagnosing problems in this area. The default value is N.
IRA_DEBUG_EVENTEXPORT	Turns on (Y) or off (N) agent event export operations, such as SNMP trap and debug trace. The default value is N.
IRA_DEBUG_PRIVATE_SITUATION	Turns on (Y) or off (N) debug trace when processing the private situations of an agent. The default value is N.
IRA_DEBUG_SERVICEAPI	Turns on (Y) or off (N) debug trace for agent service interface processing. The default value is N.
IRA_DEBUG_TRANSCON	Turns on (Y) or off (N) debug trace for agent transport conduit processing. This variable is used for diagnosing problems in this area. The default value is N.
IRA_DEBUG{EIF	Turns on (Y) or off (N) debug trace for the EIF event emitting feature in this agent.
IRA_DUMP_DATA	Used by both agents and the Tivoli Enterprise Monitoring Server for debugging. Set to Y to do a hexadecimal dump of RPC transaction content data contents into the RAS1 log. The default value is N. Can produce voluminous RAS1 message output if enabled.
IRA_EVENT_EXPORT_CHECKUSAGE_INTERVAL	Specifies the frequency with which the agent checks and calculates the autonomous operation saved event usage limit that is defined by the IRA_AUTONOMOUS_LIMIT parameter. The default value is 90 seconds; the agent enforces the minimum setting of 30 seconds.
IRA{EIF_ENABLE_LOG	Enables logging of EIF events sent for private situations directly from the Tivoli Enterprise Monitoring Agent to the agent's *.LG0 operational log if set to "Y".
IRA_EVENT_EXPORT{EIF	Enables (Y) or disables (N) the EIF event emitting feature for any private situations configured for this agent.
IRA_EVENT_EXPORT_LISTSTAT_INTERVAL	Defines the frequency with which the agent outputs collected situation statistics to the debug trace log. The default value is 900 seconds, or 15 minutes.
IRA_EVENT_EXPORT_LISTSTAT_OUTPUT	Enables (Y) or disables (N) periodic output of situation operation statistics data to the trace log. The default is N.
IRA_EVENT_EXPORT_SIT_STATS	Enables (Y) or disables (N) basic situation operation statistics data collection. The basic situation data includes situation first start time, situation first event time, situation last start time, situation last stop time, situation last true event time, situation last false event time, number of times situation recycles, number of times situation enters autonomous operation. The default value is Y.
IRA_EVENT_EXPORT_SIT_STATS_DETAIL	Enables (Y) or disables (N) details situation operation statistics data collection. The detail data collected includes 8 days' situation operation profile such as hourly true event count, hourly false event count, hourly data row count, hourly true event ratio, and hourly false event ratio. The default value is N.

Table 166. Common agent environment variables (continued)

Variable	Purpose
IRA_EVENT_EXPORT_SNMP_TRAP	Enables (Y) or disables (N) the SNMP trap emitter capability. When enabled, the SNMP trap configuration file is required and must exist for the agent to emit SNMP V1, V2, or V3 trap to configured SNMP trap managers. The default value is Y.
IRA_EVENT_EXPORT_SNMP_TRAP_CONFIG	Specifies the fully qualified SNMP trap configuration file name. The default file is located in \$CANDLE_HOME/localconfig/ <i>pc</i> / <i>pc</i> _trapcnfg.xml (member <i>pc</i> TRAPS on z/OS systems), where <i>pc</i> is the agent product code.
IRA_LOCALCONFIG_DIR	Specifies the local configuration directory path that contains locally customized configuration files such as threshold overrides, private situations, and SNMP trap configuration file. The default directory is the localconfig subdirectory of the directory specified by the CANDLE_HOME environment variable, which is the RKANDATV DD name on z/OS systems.
IRA_PRIVATE_SITUATION_CONFIG	Specifies the fully qualified autonomous Private Situation configuration file name. The default file on distributed systems is located in \$CANDLE_HOME/localconfig/ <i>pc</i> / <i>pc</i> _situations.xml, where <i>pc</i> is the agent product code. The default file on z/OS systems is the SICNFG member in the RKANDATV data set.
IRA_SERVICE_INTERFACE_DEFAULT_PAGE	<p>Instructs the agent to open the named product-specific HTML page instead of the default navigator.htm page when logging on to the agent service interface. By default, the agent looks for the product-specific file in CANDLE_HOME/localconfig on distributed systems and the RKANDATV dataset on z/OS systems. However, if the IRA_SERVICE_INTERFACE_DIR environment variable has been set, the agent looks in the directory specified by that environment variable.</p> <p>If you set IRA_SERVICE_INTERFACE_DEFAULT_PAGE (but not IRA_SERVICE_INTERFACE_DIR), you should put any product-specific HTML pages in the CANDLE_HOME/localconfig/html directory on distributed systems. Therefore, if you create myPage.htm and put it in CANDLE_HOME/localconfig/html set IRA_SERVICE_INTERFACE_DEFAULT_PAGE=/html/myPage.htm.</p>
IRA_SERVICE_INTERFACE_DIR	<p>Defines the path specification of the agent service interface HTML directory. In conjunction with the IRA_SERVICE_INTERFACE_DEFAULT_PAGE parameter, the agent constructs the file path to a specific, requested HTTP GET object. The default filepath is CANDLE_HOME/localconfig on distributed systems and the RKANDATV dataset on z/OS systems. The parameter is equivalent to the IRA_HTML_PATH parameter.</p> <p>Example: If IRA_SERVICE_INTERFACE_DIR="\mypath\private" and you enter http://localhost:1920//kuxagent/kuxagent/html/myPage.htm in your browser, myPage.htm is retrieved from \mypath\private\html\ instead of %CANDLE_HOME%\localconfig\html\.</p>
IRA_SERVICE_INTERFACE_NAME	Specifies a unique agent interface name to represent this agent. The default agent service interface name is <i>pc</i> agent, where <i>pc</i> is the application product code. In the scenario where multiple instances of the same agent are running in the system, this parameter enables customization of a unique service interface name to correspond to this agent.

Table 166. Common agent environment variables (continued)

Variable	Purpose
ITM_MANIFEST_PATH= <i>file_loc</i>	<p>The location of the self-describing agent's manifest and JAR files. Normally this parameter is set during agent installation and needs to be changed only if the files are moved to a network file system.</p> <p>File locations:</p> <ul style="list-style-type: none"> On Windows: <inst_dir>\TMAITM6\support\<pc> <inst_dir>\TMAITM6_x64\support\<pc> On Linux or UNIX: <inst_dir>/<binarch>/<pc>/support <p>where pc is the two-character product code.</p> <p>This parameter is an internal parameter that is set and used by the agent installers only. You should not change this variable.</p>
KBB_RAS1	<p>Sets the level of agent tracing:</p> <p>ERR (UNIT:KRA ST) View the state of main agent functions such as situation and report processing.</p> <p>ERR (UNIT:KRA ALL) View detailed debug messages for agent functions.</p> <p>ERR (UNIT:KHDX ST) View the state of the agent's short-term history data uploads to the Tivoli Data Warehouse.</p> <p>ERR (UNIT:KHD ALL) View detailed debugging messages about short-term history data uploads to the Tivoli Data Warehouse.</p>
KCA_CACHE_LIMIT	<p>Maximum time (in hours) that an alert is cached and seen in the Alerts view of the Agent Managent Service (AMS) workspace.</p> <p>Possible values: Greater than 0.</p> <p>Default value: 24.</p>
KCA_CAP_DIR	<p>Directories looking for CAP files.</p> <p>Possible values: Any valid PATH.</p> <p>Default value:</p> <ul style="list-style-type: none"> On Windows systems: %CANDLE_HOME%\TMAITM6_x64\CAP;@BinPath@\CAP;%ALLUSERSPROFILE%\Application Data\IBM\CAP On Linux or UNIX: \$CANDLEHOME\$/config/CAP:/opt/IBM/CAP
KCA_CAP_DISCOVERY_INTERVAL	<p>Time interval (in seconds) for discovering new or changed CAP files.</p> <p>Possible values: Greater than 0.</p> <p>Default value: 30.</p>
KCA_CMD_TIMEOUT	<p>Timeout (in seconds) for external commands created by PAS.</p> <p>Possible values: Greater than or equal to 90.</p> <p>Default value: 90.</p>

Table 166. Common agent environment variables (continued)

Variable	Purpose
KCA_DISCOVERY_INTERVAL	<p>How frequently the watchdog polls for new, running instances of agents managed by Agent Management Services.</p> <p>Possible values: Greater than 0.</p> <p>Default value:</p> <ul style="list-style-type: none"> On Windows systems: 30 On Linux or UNIX: 120
KCA_IP_DIR	<p>Only for Linux and UNIX; directory for creating sockets to accept external commands.</p> <p>Possible values: Any valid existing PATH.</p> <p>Default value: \$CANDLEHOME\$/ \$BINARCH\$/ \$PRODUCTCODE\$/bin/pasipc</p> <p>Note:</p> <ol style="list-style-type: none"> A complete list of the platforms (\$BINARCH) and codes can be found in <ITMHOME>/registry/archdsc.tbl. Some systems can have a problem managing sockets with long path names. When installing the agent on a CANDLEHOME path that is longer than 50 characters you should set KCA_IP_DIR to a value shorter than the default value. For example KCA_IP_DIR=/tmp/pasipc.
KCA_ITM_DISCOVERY_INTERVAL	<p>Time interval (in seconds) for discovering new unmanaged agent instances of IBM Tivoli Monitoring agents (cinfo command).</p> <p>Possible values: Greater than KCA_DISCOVERY_INTERVAL.</p> <p>Default value: 600.</p>
KCA_MAX_RETRIES_ON_PIPE	<p>Maximum number of consecutive failed attempts to check the availability of the OS agent through the PAS internal socket (pipe). When KCA_MAX_RETRIES_ON_PIPE is reached, the OS agent is restarted. An undefined variable, or less than 0, means an infinite number of retries (no restart).</p> <p>Possible values: Greater than 5.</p> <p>Default value: Not defined.</p>
KCAWD_WRITE_DIR	<p>Writable directory for kcawd (WatchDog agent).</p> <p>Possible values: Any valid directory.</p> <p>Default value: The installation directory of the OS agent binary.</p>
KHD_HISTSIZE_EVAL_INTERVAL	<p>Controls how often (in seconds) the Warehouse Proxy Agent client code checks the size of the historical files. For more information, see the <i>IBM Tivoli Monitoring Administrator's Guide</i>.</p>
KHD_HISTRETENTION	<p>Specifies the default retention period in hours for the short-term history files (default is 24 hours). This value can be used to reduce the amount of data kept on disk after a successful upload to the warehouse is performed.</p>
KHD_STATUSTIMEOUT	<p>The time in seconds set by default to 900s = 15 minutes. An export request on the application agent is resent if a status is not received from the Warehouse Proxy agent before the timeout expires. It is advisable not to change this variable.</p>

Table 166. Common agent environment variables (continued)

Variable	Purpose
KHD_TOTAL_HIST_MAXSIZE	Controls the threshold that the Warehouse Proxy Agent client code uses when checking the size of the historical files. When their combined size is over the specified value in MB, writing to them is suspended until the threshold is no longer exceeded. A value of 0 disables the periodic checking. For more information, see the <i>IBM Tivoli Monitoring Administrator's Guide</i> .
TIMEOUT	Specifies the time in seconds that Agent Deployment tool has to complete a task. If the tool does not complete in the task in the time specified by the TIMEOUT value, the task is terminated. The default value is 600 seconds.
TEMA_SDA	N disables the self-describing agent capability at the agent, whereas Y enables it. A value of N blocks the Tivoli Enterprise Monitoring Server from retrieving any product support files from this agent and provides you with control on a per agent basis without stopping the self-describing agent feature on the Tivoli Enterprise Monitoring Server for other products. Note: Do not specify YES or NO; instead, always specify Y or N.
TEMA_SDA_ACK_WAIT	Time interval in seconds for the agent to wait on a response from the Tivoli Enterprise Monitoring Server for the completion status of an actual self-describing agent installation. A valid value is a number in seconds greater than or equal to 60. The default value is 300 seconds. This parameter is an internal parameter that you might use, in a rare situation, to fine-tune the self-describing agent capability. In general, accept the default value of this variable.
TEMA_SDA_RETRY_WAIT	Time interval in seconds for the agent to wait before attempting another self-describing agent installation request. Not all failed self-describing agent registration or installation responses from the Tivoli Enterprise Monitoring Server are retried. A valid value is a number in seconds greater than or equal to 60. The default value is 600 seconds. This parameter is an internal parameter that you might use, in a rare situation, to fine-tune the self-describing agent capability. In general, accept the default value of this variable.
TEMA_SDA_MAX_ATTEMPTS	The maximum number of times the agent tries to register a request to initiate a self-describing agent product installation with the Tivoli Enterprise Monitoring Server. The agent will continue to try its self-describing agent registration until the Tivoli Enterprise Monitoring Server self-describing agent manager has completed the installation of this agent's product support. The minimum value is 1. The agent must always try its self-describing agent registration once (unless the TEMA_SDA value is N). This maximum attempt count should only be modified with direction from IBM Software Support. If the value is set too low, your self-describing agent product installation might not complete as expected. The agent must be able to retry its self-describing agent registration until the product installation has completed. In the event of an agent attempting excessive self-describing agent product install retry attempts, it might generate excessive or unwanted audit or error messages that you want to stop. Decreasing this number can help in such a case. This parameter is an internal parameter that you might use, in a rare situation, to fine-tune the self-describing agent capability. In general, accept the default value of this variable.

Operating system agent environment variables

This section lists IBM Tivoli Monitoring OS agent environment variables that you can customize

Windows OS monitoring agent

The table in this section lists the Windows OS monitoring agent environment variables.

Table 167. Windows OS monitoring agent environment variables

Variable	Description	Possible values	Default value
KNT_CPUSTAT_SAMPLE_SECS	To customize attribute values for the total CPU metrics. For more information, see the <i>Windows OS Agent User's Guide</i> .	0, or greater than or equal to 5.	30 seconds
KNT_PROCESS_SAMPLE_SECS	To customize attribute values for the CPU metrics per process. For more information, see the <i>Windows OS Agent User's Guide</i> .	0, or greater than or equal to 5.	60 seconds
NT_EXCLUDE_PERF_OBJS	To explicitly exclude collection of metrics from the specified Performance Objects.	List of comma-separated strings.	Blank
NT_EXCLUDE_UNMAPPED_DISKS	To avoid reporting of unmapped hard disks among Logical Disks (having instance name not mapped to a letter yet).	0 or 1.	0
NT_LOG_DUPLICATE	To enable reporting of the Duplicate Record Count for duplicated events in the event log.	0 or 1.	0
NT_LOG_MAX_EVTS	To define the max number of past events to be searched in the Event Log. 0 means that the missed events feature is disabled.	0, or greater than or equal to 1.	0
NT_LOG_MAX_TIME	To define the maximum interval, in seconds, of past events to be searched in the Event Log. 0 means that the missed events feature is disabled.	0, or greater than or equal to 1.	0
NT_LOG_THROTTLE	To drop the specified number of duplicate events every read cycle of the event log. 0 means throttling is disabled.	0, or greater than or equal to 1.	0
NT_PERFMON_MEMORY_CHECK	To disable the exclusion of Performance Objects that are leaking memory, as a result of the check performed by the agent at initialization time.	0 or 1.	1
REVERSE_LOOKUP_ACCEPTED_FAILURES	To avoid long Network Port attribute group response times due to reverse lookup failures when retrieving Local and Remote Names from corresponding IP addresses. 0 means that reverse lookup is disabled.	0, or greater than or equal to 1.	30

Linux OS monitoring agent

The table in this section lists the Linux OS monitoring agent environment variables.

Table 168. Linux OS monitoring agent environment variables

Variable	Description	Possible values	Default value
KBB_NFS_TIMEOUT	To adjust the timeout on NFS file systems monitoring (in seconds).	1 - 30.	2
KBB_SHOW_NFS	Specifies whether NFS monitoring is enabled. The default value is false on Linux and UNIX systems.	[true false], not case sensitive.	false, not case-sensitive
KDEBE_FIPS_MODE_ENABLED	To request GSKit compliance to FIPS 140-2 standards when computing checksums for files in the File Information attribute group.	[yes no], not case-sensitive.	no, not case-sensitive

Table 168. Linux OS monitoring agent environment variables (continued)

Variable	Description	Possible values	Default value
KLZ_CPUSTAT_SAMPLE_SECS	To tune the sampling interval for the CPU attribute group (in seconds).	Greater than or equal to 5.	30
KLZ_DISK_SAMPLE_HRS	To tune the sampling interval for the Disk Usage Trends attribute group (in seconds).	Greater than or equal to 1.	3600
KLZ_IOSTAT_SAMPLE_SECS	To tune the sampling interval for the Disk IO attribute group (in seconds).	Greater than or equal to 5.	30
KLZ_NETSTAT_SAMPLE_SECS	To tune the sampling interval for the Network attribute group (in seconds).	Greater than or equal to 5.	30
KLZ_PINGHOSTLIST	To set the full qualified path of the file containing the list of servers for the Host Availability attribute group.	A valid UNIX pathname.	Blank
KLZ_PROCESS_SAMPLE_SECS	To tune the sampling interval for the Process Instant Busy CPU (Percent) attribute in the Process group (in seconds).	0, or greater than or equal to 5.	60
KLZ_PROCNORM_TEST	To enable Process short-term CPU average attributes also for Queries.	[y n], not case-sensitive.	n, not case-sensitive
KLZ_SETLPARVMID	To set the CTIRA_HOSTNAME parameter for the agent as Lpname.VMname:LZ (Linux 390 systems only).	[y n], not case-sensitive.	n, not case-sensitive
KLZ_SKIP_DISABLED_DISKS	To avoid the OS agent reporting on SAN Physical Disks that are not available to the system.	[TRUE true True] or false, not case-sensitive.	false, not case-sensitive
KLZ_SWAPTREND_SAMPLE_HRS	To tune the sampling interval for the Swap Rate attribute group (in seconds).	Greater than or equal to 1.	3600
KLZ_SYSSTAT_SAMPLE_SECS	To tune the sampling interval for the System Statistics attribute group (in seconds).	Greater than or equal to 5.	30
KLZ_TCPSTAT_SAMPLE_SECS	To tune the sampling interval for the TCP Statistics attribute group.	Greater than or equal to 5.	30
UID_USERNAME_REFRESH_TIME	To tune the refresh interval of uid<->username and gid<->groupname mapping tables (in minutes).	Greater than or equal to 1.	60

UNIX OS monitoring agent

The table in this section lists the UNIX OS monitoring agent environment variables.

Table 169. UNIX OS monitoring agent environment variables

Variable	Description	Possible values	Default value
AGENT_TIMING	To enable agent internal timers.	Y	Y
EXCLUDE_LEGACY_DISKS	To exclude monitoring of disks that follow the HPUNIX legacy naming convention.	[TRUE true True] or false, not case-sensitive.	false, not case-sensitive
KBB_EXCLUDE_SLOW_ISATTY	To check the response time of the isatty() system call and eventually skip it to avoid agent hangs.	0 or 1.	0
KBB_SHOW_CUSTOMFS	To enable monitoring of a specific filesystem type.	Between 0 and 30.	blank
KBB_SHOW_NFS	Specifies whether NFS monitoring is enabled. The default value is false on Linux and false on UNIX.	[true false], not case-sensitive.	false, not case-sensitive

Table 169. UNIX OS monitoring agent environment variables (continued)

Variable	Description	Possible values	Default value
KUX_CPUSTAT_SAMPLE_SECS	To customize attribute values for the total CPU metrics. For more information, see the <i>UNIX OS Agent User's Guide</i> .	0, or greater than or equal to 5.	30 seconds
KUX_PROCESS_SAMPLE_SECS	To customize attribute values for the CPU metrics per process. For more information, see the <i>UNIX OS Agent User's Guide</i> .	0, or greater than or equal to 5.	60 seconds
KUX_AIXDP	To disable the AIX data provider metrics collection (new attributes inherited from the AIX Premium Agent). Note: The minimum AIX requirements to collect metrics from UNIX OS monitoring agent and Premium Monitoring Agent for AIX convergence are: <ul style="list-style-type: none"> AIX61F = AIX 6.1 TL5 	[true false], not case-sensitive.	true, not case-sensitive
KUX_DEFINED_USERS	To enable the Defined Users group on AIX systems.	True or false, not case-sensitive.	false, not case-sensitive
KUX_GETARGS_DELAY	To adjust the interval between subsequent attempts to get Processes command-line arguments (milliseconds - AIX systems only).	Between 0 and 1000.	10
KUX_GETPROCS_DELAY	To adjust the sampling interval for Process CPU percent utilization (milliseconds - AIX only).	0, or greater than or equal to 1.	2000
KUX_PINGHOSTLIST	To set the full qualified path of the file containing the list of servers for the UNIX Ping attribute group.	A valid UNIX pathname.	blank
KUX_PRCTL_OFF	To disable collecting CPUSHARES and SHAREPCT to avoid crashes in the prctl command due to a Solaris Zones bug.	[TRUE true True] or false, not case-sensitive.	false, not case-sensitive
KUX_SKIP_DISABLED_DISKS	To avoid the OS agent reporting on SAN Physical Disks that are not actually available to the system.	[TRUE true True] or false, not case-sensitive.	false, not case-sensitive
KUX_SKIP_HP_UTMPX	To skip counting the number of user sessions for the All Users group on HP-UX.	True or false, not case-sensitive.	false, not case-sensitive
KUX_SKIP_TTY	To improve process monitoring performances by skipping TTY names resolution.	[TRUE true True] or false, not case-sensitive.	false, not case-sensitive
KUX_SKIP_UTF8CONV_PROCESS	To reduce OS agent high CPU consumption when converting strings to UTF8 for large number of Processes.	[TRUE true True] or false, not case-sensitive.	false, not case-sensitive
KUX_TCPSTAT_SAMPLE_SECS	To tune the sampling interval for the TCP Statistics attribute group (in seconds).	Greater than or equal to 5.	30
KBB_NFS_TIMEOUT	To tune the timeout on NFS filesystems monitoring (in seconds).	Between 1 and 30.	2
KDEBE_FIPS_MODE_ENABLED	To request GSKit compliance to FIPS 140-2 standards when computing checksums for files in the File Information attribute group.	[yes no], not case-sensitive.	no, not case-sensitive
MAX_NUMBER_OF_DISKS_HP-UX	To define an upper limit on the number of disks that the agent can monitor (HP-UX only).	Between 0 and 30000.	3500

Table 169. UNIX OS monitoring agent environment variables (continued)

Variable	Description	Possible values	Default value
NETWORK_INTERFACE_REFRESH_TIME_HPUX	To tune the frequency of full refresh of network interfaces data (in hours - HPUX only).	Between 1 and 168.	0
TTY_REFRESH_TIME	To tune the refresh interval of tty<->ttyname mapping table (in minutes).	Greater than or equal to 1.	60
UID_USERNAME_REFRESH_TIME	To tune the refresh interval of uid<->username and gid<->groupname mapping tables (in minutes).	Greater than or equal to 1.	60

Appendix F. Maintaining the EIB on Linux or UNIX

To ensure the effective operation of your monitoring server, back up your Enterprise Information Base (EIB) tables as part of your routine maintenance.

Enterprise Information Base (EIB) tables

The EIB contains the attributes and other data that define the agents to the server. The following files, which are stored in the *install_dir*/tables/eib directory, compose the EIB.

Table 170. EIB Files

*.db Files		*.idx Files	
qa1cacts.db	qa1daggr.db	qa1cacts.idx	qa1daggr.idx
qa1cckpt.db	qa1dcct.db	qa1cckpt.idx	qa1dcct.idx
qa1ccobj.db	qa1dcct2.db	qa1ccobj.idx	qa1dcct2.idx
qa1ccomm.db	qa1dmobj.db	qa1ccomm.idx	qa1dmobj.idx
qa1ceibl.db	qa1dmtmp.db	qa1ceibl.idx	qa1dmtmp.idx
qa1chost.db	qa1dobja.db	qa1chost.idx	qa1dobja.idx
qa1ciobj.db	qa1dpcyf.db	qa1ciobj.idx	qa1dpcyf.idx
qa1cmcfg.db	qa1drnke.db	qa1cmcfg.idx	qa1drnke.idx
qa1cnodl.db	qa1drnkg.db	qa1cnodl.idx	qa1drnkg.idx
qa1cplat.db	qa1dsnos.db	qa1cplat.idx	qa1dsnos.idx
qa1cpset.db	qa1dspst.db	qa1cpset.idx	qa1dspst.idx
qa1cruld.db	qa1dstms.db	qa1cruld.idx	qa1dstms.idx
qa1csitf.db	qa1dstsa.db	qa1csitf.idx	qa1dstsa.idx
qa1csmni.db	qa1dstua.db	qa1csmni.idx	qa1dstua.idx
qa1cstsc.db	qa1dswrs.db	qa1cstsc.idx	qa1dswrs.idx
qa1cstsh.db	qa1dswus.db	qa1cstsh.idx	qa1dswus.idx
qa1cthre.db	qa1dwgrp.db	qa1cthre.idx	qa1dwgrp.idx
qa1dactp.db	qa1dwork.db	qa1dactp.idx	qa1dwork.idx

Appendix G. Securing your IBM Tivoli Monitoring installation on Linux or UNIX

The **secureMain** utility helps you bring the monitoring environment into compliance with the security standards of your company.

On UNIX and Linux operating systems, the product installation process creates the majority of directories and files with **world write** permissions. This configuration creates a security situation that is not acceptable in many enterprises. Run the **secureMain** utility on all installations, especially those installations that include the UNIX OS Agent, to prevent privilege escalation.

Note: You do not need to be logged in as a root user to run this utility, but you are prompted for the root password when it is required.

Usage

Before you run the **secureMain** utility, ensure that you understand the syntax variables.

The **secureMain** commands use the following syntax:

```
secureMain [-h install_dir] [-g common_group] [-t type_code] lock
secureMain [-h install_dir] [-g common_group] unlock
```

where variables are defined as follows:

- ***install_dir*** is the directory path for the IBM Tivoli Monitoring installation. If this parameter is not supplied, the script attempts to determine the installation directory.
- ***common_group*** is a group ID common to all of the user IDs that are used to run components in this installation. The user ID that is used to perform the installation must also be a member of the group ID specified. The only exception is that the **root** ID is not required to be a member of the group ID specified.
- ***type_code*** is a component code belonging to an installed component. You can specify multiple **-t** options to create a list of component codes to be processed.

If **secureMain** is invoked with no parameters, the usage text is displayed.

secureMain lock is used to tighten permissions in an IBM Tivoli Monitoring 6.1 installation. It should be run after installing or configuring components.

When **secureMain lock** is invoked with no other parameters, the permissions are tightened generally to 755. However, a number of directories and some files are still left with **world write** permissions. When certain components which are commonly run using multiple user IDs are present in the installation, many more files have **world write** permissions.

When **secureMain lock** is invoked with the **-g *common_group*** parameter, the permissions are tightened generally to 775 and the directories and files have their group owner changed to ***common_group*** specified. There are no directories or files left with **world write** permissions. Even when certain components which are commonly run using multiple user IDs are present in the installation, no files will have **world write** permissions. Additionally, the ***common_group*** value specified is written to a file and is used for all future **secureMain lock** invocations in this installation, unless the **-g** option is specified and the ***common_group*** is different from the previous value.

When **secureMain lock** is invoked with the **-t *type_code*** parameter, sections of the installation might be skipped when tightening permissions. Common directories, like **bin**, **config**, **registry**, and **logs**, and the

files in them are always processed. Only directories and files specific to the specified **type_code** components are processed. The other component directory trees are skipped.

secureMain unlock is used to loosen permissions in an IBM Tivoli Monitoring installation. **secureMain unlock** is normally not necessary, but can be run if desired. It should be run before installing or configuring components.

secureMain unlock does not return the installation to the permission state that it was in before running **secureMain lock**. It only processes the common directories, like bin, config, registry, and logs, and the files in them.

Example

The following example locks the installation using the common group **itmgroup**:

```
secureMain -g itmgroup lock
```

The following example locks the base and mq component directories using the common group **itmgroup**:

```
secureMain -g itmgroup -t mq lock
```

Scenario with secureMain

The following scenario illustrates the use of **secureMain**:

1. Complete the following operations using **root** authorization:
 - a. Install OS Agent.
 - b. Configure OS Agent.
 - c. List files with **world write** permissions, using the following command: `find . -perm -o+w -ls`
 - d. Run the following command: `secureMain -g itmgroup -t ux lock`
 - e. Install the 32-bit Enterprise Svcs UI to get the 32-bit framework.
 - f. Install the MQ agent.
 - g. Run the following command: `secureMain -g itmgroup -t mq lock`
 - h. List files with **world write** permissions, using the following command: `find . -perm -o+w -ls`
 - i. Start the OS agent.
2. Complete the following operations using **mquser** authorization:
 - a. Start the MQ agent for a queue manager.
 - b. Start the MQ agent for a second queue manager.
 - c. Stop the MQ agent for the first queue manager.
 - d. Stop the MQ agent for the second queue manager.
3. Complete the following operations using **root** authorization:
 - a. Stop the OS Agent.
 - b. List files with **world write** permissions, using the following command: `find . -perm -o+w -ls`

Securing your IBM Tivoli Monitoring environment at installation

The command `./install.sh` contains the optional `-k` group parameter. The `-k` parameter runs the **secureMain** utility at the end of the installation process.

Example

```
./install.sh -h /opt/IBM/ITM -k itmgroup
```

If you do not specify the `-k` parameter when you run `./install.sh`, you will be asked at the end of the installation if you want to secure your IBM Tivoli Monitoring environment. If your IBM Tivoli Monitoring

environment is already secured this question is skipped. When you install a monitoring agent into an already secured environment, **secureMain** will automatically execute at the end of the installation.

“Installing into an existing installation” on page 137

Installing components or agents in an existing CANDLEHOME or installation directory is supported as long as the user ID used to run the installation is always the same.

Appendix H. Uninstalling IBM Tivoli Monitoring

This appendix provides information about uninstalling IBM Tivoli Monitoring.

Use the following steps to uninstall IBM Tivoli Monitoring:

- “Uninstalling the entire IBM Tivoli Monitoring environment”
- “Uninstalling an individual IBM Tivoli Monitoring agent or component” on page 934
- “Uninstalling components and agents silently” on page 938
- “Uninstalling the event synchronization component” on page 940
- “Uninstalling the Tivoli Authorization Policy Server or tivcmd Command Line Interface for Authorization Policy using the graphical user interface” on page 947
- “Uninstalling the IBM Infrastructure Management Dashboards for Servers using the graphical user interface” on page 947

Uninstalling the entire IBM Tivoli Monitoring environment

This section contains information for uninstalling the entire IBM Tivoli Monitoring base environment from Windows and Linux/UNIX systems.

Note: This procedure does not uninstall the components for a dashboard environment such as IBM Infrastructure Management Dashboard for Servers, Tivoli Authorization Policy Server, and tivcmd Command Line Interface for Authorization Policy.

Use the following procedures to remove the entire IBM Tivoli Monitoring environment from your computer.

- “Uninstalling the environment on Windows”
- “Uninstalling the environment on Linux or UNIX” on page 933

If you want to remove just one component such as an agent, see “Uninstalling an individual IBM Tivoli Monitoring agent or component” on page 934.

Note: If you plan to reinstall IBM Tivoli Monitoring into a different directory than the one used for this installation, you must stop and restart this computer before reinstalling IBM Tivoli Monitoring.

Uninstalling the environment on Windows

Use the steps in this procedure to uninstall IBM Tivoli Monitoring from a Windows computer.

Procedure

1. From the desktop, open the Control Panel.
2. Click **Add/Remove Programs**.
3. Select **IBM Tivoli Monitoring** and click **Change/Remove**. The following window is displayed.

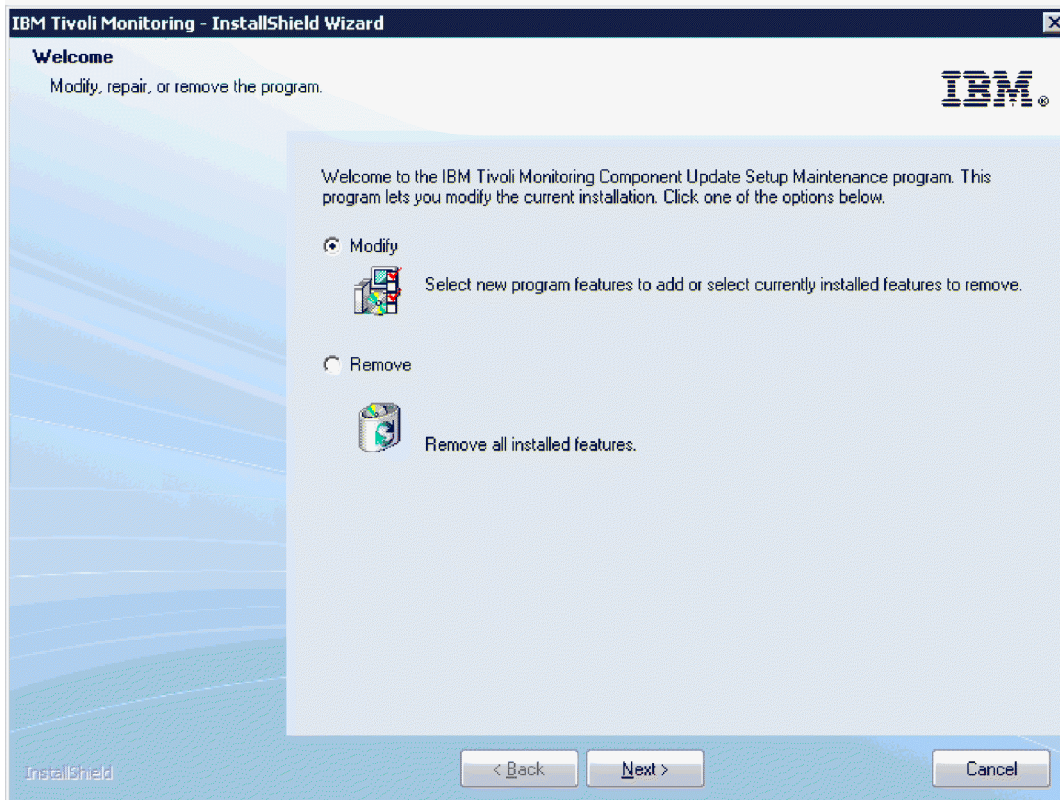


Figure 168. Uninstalling IBM Tivoli Monitoring

4. Select **Remove** and click **Next**.
The following window is displayed.




Figure 169. Confirming the uninstallation

5. Click **OK**.
The following progress window is displayed.




Figure 170. Stopping Tivoli components before uninstallation

After Tivoli Enterprise services have stopped, you are asked if you want to remove the Tivoli Enterprise Portal database.




Figure 171. Removing the portal database

6. Click **Yes**.

The following window is displayed, requesting information required to remove the database:




Figure 172. Database information

7. Type the password for the database administrator in the **Admin Password** field and click **OK**.
The following progress window is displayed.




Figure 173. Uninstallation progress window

A pop-up window, indicating that GSKit is being uninstalled, is displayed.




Figure 174. GSKit uninstallation

After GSKit is uninstalled, the following window is displayed:

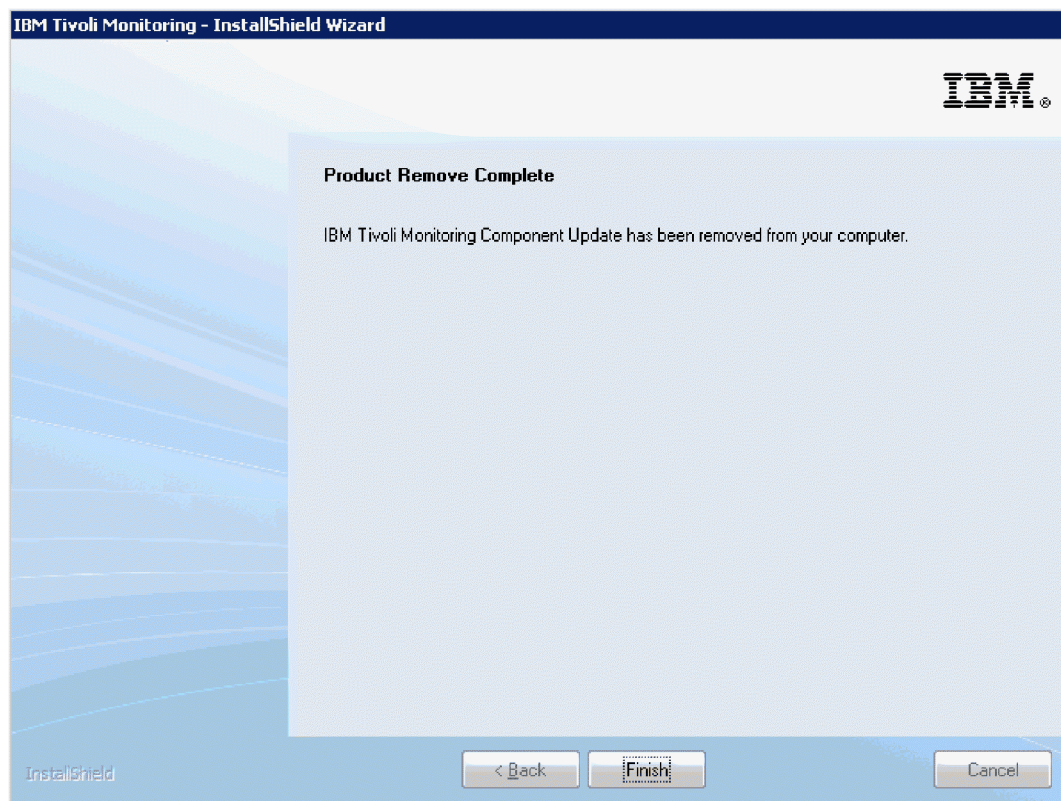


Figure 175. Successful uninstallation

8. Click **Finish**.

Uninstalling the environment on Linux or UNIX

Use the steps in this procedure to uninstall IBM Tivoli Monitoring from a UNIX computer.

Procedure

1. From a command prompt, run the following command to change to the appropriate `/bin` directory:
`cd install_dir/bin`

where *install_dir* is the path for the home directory for IBM Tivoli Monitoring.

2. Run the following command:
`./uninstall.sh`

A numbered list of product codes, architecture codes, version and release numbers, and product titles is displayed for all installed products.

3. Type the number for the installed product that you want to uninstall. Repeat this step for each additional installed product you want to uninstall.
4. After you have removed all installed components, you are asked if you want to remove the installation directory. Type `y` and press Enter.

What to do next

You can also run the following command to remove all installed components from the command-line:
`./uninstall.sh REMOVE EVERYTHING`

After the command completes, you can manually remove the IBM Tivoli Monitoring installation directory.

Note: If for any reason, the UNIX uninstallation is not successful, run the following command to remove all IBM Tivoli Monitoring directories:

```
rm -r install_dir
```

This uninstallation program does not delete the database created for Tivoli Enterprise Portal on a Linux portal server. If you want to delete that database, you must remove it manually. See the documentation for your database software for information about deleting a database.

Uninstalling an individual IBM Tivoli Monitoring agent or component

Use the procedures in this section to remove an agent or other individual IBM Tivoli Monitoring component from your computer.

- “Uninstalling a component on Windows”
- “Uninstalling a component on Linux or UNIX” on page 935
- “Uninstalling OMEGAMON Monitoring Agents” on page 935
- “Removing an agent through the Tivoli Enterprise Portal” on page 937

Note: If you plan to reinstall a IBM Tivoli Monitoring component into a different directory than the one used for this installation, you must stop and restart this computer before reinstalling the IBM Tivoli Monitoring component.

Uninstalling a component on Windows

Use the steps in this section to remove a component on a Windows computer.

About this task

You can uninstall a single agent or the entire agent bundle (such as IBM Tivoli Monitoring for Databases). You cannot uninstall the application support files laid down for a Windows Tivoli Enterprise Monitoring Server without uninstalling the monitoring server.

Procedure

1. From the desktop, open the Control Panel.
2. Click **Add/Remove Programs**.
3. Complete one of the following:
 - To uninstall a single IBM Tivoli Monitoring component, such as the portal server or portal client (but not all components), select **IBM Tivoli Monitoring**.
 - To uninstall an agent bundle or a specific agent, select the agent bundle.
4. Click **Change/Remove**.
5. Complete one of the following steps:
 - To uninstall a specific agent or component, select **Modify**.
 - To uninstall the entire agent bundle, select **Remove**.
6. Click **Next**.
7. If you are uninstalling an agent bundle, click **OK** to confirm the uninstallation.
8. If you are uninstalling an agent or component, complete the following steps:
 - a. For an agent, expand **Tivoli Enterprise Monitoring Agents** and select the agent you want to uninstall.
 - b. For a component, select the component (such as **Tivoli Enterprise Portal Desktop Client**).
 - c. For Tivoli Performance Analyzer, clear the Tivoli Performance Analyzer check box in all sections (such as Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Tivoli Enterprise Portal Desktop Client, and so on).

- d. Click **Next**.
- e. Click **Next** on the confirmation screen.
- f. Depending on the remaining components on your computer, there might be a series of configuration panels. Click **Next** on each of these panels.

Note: When you are uninstalling the Tivoli Enterprise Portal Server, the installer gives you the option to remove the portal server database. If there are other databases created by Tivoli Monitoring in this or previous version on the computer, the installer gives you the option to remove them as well.

9. Click **Finish** to complete the uninstallation.

Uninstalling a component on Linux or UNIX

Use the steps in this section to remove a component on a UNIX computer.

About this task

You can uninstall a single agent or the entire agent bundle (such as IBM Tivoli Monitoring for Databases).

Procedure

1. From a command prompt, run the following command to change to the appropriate /bin directory:

```
cd install_dir/bin
```

where *install_dir* is the path for the home directory for IBM Tivoli Monitoring.

2. Run the following command:

```
./uninstall.sh
```

A numbered list of product codes, architecture codes, version and release numbers, and product titles is displayed for all installed products.

3. Type the number for the agent or component that you want to uninstall. Repeat this step for each additional installed product you want to uninstall.

Note: When you are uninstalling the Tivoli Enterprise Portal Server, the installer gives you the option to remove the portal server database. If there are other databases created by Tivoli Monitoring in this or previous version on the computer, the installer gives you the option to remove them as well.

Uninstalling OMEGAMON Monitoring Agents

Use the steps in this section to remove OMEGAMON agents from a computer.

Before you begin

Table 171 on page 936 and Table 172 on page 936 list the agents by internal code, release, and descriptive name.

Procedure

1. Launch Manage Candle Services (350 or 360) or Manage Tivoli Enterprise Monitoring Services.
2. Use the **Description** and **Release** columns to locate the agent service name.
3. Stop the service by right-clicking the name and clicking **Stop**.
4. Take note of any task or subsystem names that are listed in the column for your agent. Usually this column lists **Primary** unless your agent supports instances. If your agent supports instances, record these names for later use.
5. Unconfigure the agent by right-clicking the name and clicking **Advanced** → **Unconfigure**. The **Configured** column changes from Yes to No. Continue to unconfigure all instances found in step 4.

6. Open the Windows Explorer and navigate to the installation directory for OMEGAMON 350 or 360 products and IBM Tivoli Monitoring. The default directories are C:\Candle Candle OMEGAMON and C:\IBM\ITM for IBM Tivoli Monitoring. Then navigate to the CMA directory.
7. Delete files K??ENV (Task/SubSystem name Primary) and any instances shown as K??ENV_INSTANCENAME (Task/SubSystem name from step 4 on page 935).
8. Delete any PC*.EXE or PC*.DLL files for the product. PC is the product internal identifier three-character code from the tables.
9. Exit Manage Candle Services or Manage Tivoli Enterprise Monitoring Services, and launch it again. The agent and all instances should not be shown under the **Service/Application** column.

What to do next

You can also use this procedure to remove IBM Tivoli Monitoring agents if you use the TMAITM6 directory instead of the CMA directory in step 6. All of the other steps do not change.

Table 171. Candle OMEGAMON Release 04R1

Internal identifier	Release	Description
K3Z	400	Windows Server Active Directories Monitoring Agent
KA2	120	Alert Adapter for AF/Remote
KA4	300	Monitoring Agent for OS/400
KBL	320	CASP Directory Server Monitoring Agent
KBR	320	CASP Exchange Connector Monitoring Agent
KEZ	251	eBA Solutions Monitoring Agent
KIC	100	WebSphere InterChange Server Monitoring Agent
KIE	100	WebSphere InterChange Server Data Source
KMA	201	Alert Adapter for Remedy ARS
KMC	360	WebSphere MQ Configuration Agent
KMQ	360	WebSphere MQ Monitoring Agent
KNW	300	NetWare Monitoring Agent
KOQ	301	Microsoft SQL Server Monitoring Agent
KOR	301	Oracle Monitoring Agent
KOY	300	Sybase Monitoring Agent
KPT	201	Alert Adapter for Peregrine Service Center
KQI	120	WebSphere Integration Brokers Monitoring Agent
KSA	301	R/3 Monitoring Agent
KTX	300	Tuxedo Monitoring Agent
KUD	400	DB2 Universal Database Monitoring Agent
KWE	130	WebSphere Application Server Monitoring Agent
KWL	100	BEA WebLogic Server Monitoring Agent
KWN	100	Windows Management Web Service

Table 172. Candle OMEGAMON Release BIV110

Internal identifier	Release	Description
KIC	110	WebSphere InterChange Server Monitoring Agent
KIE	110	WebSphere InterChange Server Data Source

Table 172. Candle OMEGAMON Release BIV110 (continued)

Internal identifier	Release	Description
KMC	370	WebSphere MQ Configuration Agent
KMQ	370	WebSphere MQ Agent
KQI	130	WebSphere Integration Brokers Monitoring Agent

Removing an agent through the Tivoli Enterprise Portal

You can also uninstall non-OS monitoring agents from the Tivoli Enterprise Portal by stopping the agent and removing its configuration settings.

Before you begin

After you have removed the agent from the enterprise, you can completely uninstall the agent from the managed system. When you remove an agent, it is removed from any managed system groups to which it is assigned, any situation or policy distribution lists it was on, and any custom Navigator view items to which it was assigned.

Note: You cannot use the Tivoli Enterprise Portal to remove or uninstall OS agents.

Note: If the Manage Tivoli Enterprise Monitoring Services utility is running when you uninstall the agent, it is shut down automatically by the uninstallation process.

About this task

Use the following steps to remove and optionally uninstall an agent:

Procedure

1. In the Tivoli Enterprise Portal, right-click the agent Navigator item and click **Remove**.
2. Click **Yes** when you are asked to confirm the removal of the agent.
3. When you are asked to confirm that you want to permanently uninstall the agent, click **Yes** to uninstall or **No** to leave the agent installed on the computer.

Uninstalling the Warehouse Proxy

When you uninstall the Warehouse Proxy, the warehouse database is not removed and historical situations on the agent are not stopped.

About this task

Before you uninstall, complete the following steps to uninstall the warehouse database and historical situations on the agent:

Procedure

1. Stop the historical situations.
2. Drop the warehouse database.
3. Remove the ODBC data source.
4. Remove the Windows user, ITMUser, that was created to connect to a DB2 for Linux, UNIX, and Windows database.

Removing the ODBC data source connection

When you uninstall IBM Tivoli Monitoring, the ODBC data source created for the Warehouse Proxy Agent is not removed automatically, which can cause problems when you reinstall IBM Tivoli Monitoring. To prevent these problems, manually remove the ODBC data source after you uninstall IBM Tivoli Monitoring.

For example, to remove the DB2 for Linux, UNIX, and Windows data source from the DB2 command-line, run the following command:

```
DB2 UNCATALOG SYSTEM ODBC DATA SOURCE datasource_name
```

If you are using a Microsoft SQL database or an Oracle database, use the Windows ODBC Data Source Administrator utility to remove the ODBC data source.

Uninstalling components and agents silently

Use the following procedures to uninstall components and agents silently.

- “Performing a silent uninstallation on a Windows computer”
- “Performing a silent uninstallation on a Linux or UNIX computer” on page 939

Performing a silent uninstallation on a Windows computer

Like silent installation, silent uninstallation on Windows uses a response file. Sample response files are shipped with IBM Tivoli Monitoring components and all monitoring agents that use the IBM Tivoli Monitoring services.

The sample files can be found in one of the following locations:

- On the product installation media
- After installation, in the *install_dirsamples* directory (if present)

Note: The name of the response file varies by installation media and release. The file is typically found on the media under the WINDOWS directory. For Windows on Itanium, there is a separate WIA64 directory with an Itanium specific response file. If IBM Tivoli Monitoring was used to generate the response file for a configured agent, the response file is located where it was generated (*install_dir*\response is the default) and the name of the file is *silent_install_pc.txt*. See “Automatically creating agent response files on Windows” on page 842 for more information.

Complete the following steps to edit the response file as appropriate for your uninstallation:

1. Locate the appropriate file. For example:

silent_server.txt

For a server image

silent_agent.txt

For an agent image

silent_WIA64.txt

For an agent image for 64-bit Windows Itanium

2. Copy this file to a temporary directory on your system.

3. Open your copy of the file in a text editor.

4. Change the parameters as appropriate for the uninstallation:

- To remove all components and the installation directory, uncomment (by removing the semicolon) the following line in the ACTION TYPE section:

```
;REMOVEALL=Yes
```

- To select a component to uninstall, uncomment the following line in the ACTION TYPE section

```
;UNINSTALLSELECTED=Yes
```

and uncomment the component or components to be removed in the FEATURE section. For example:

```
*****
;
;          TIVOLI ENTERPRISE MONITORING AGENT
;          TEMA INSTALLATION SECTION
;
; Any Feature selected that ends in CMA will cause the TEMA Framework
; and specific Agent to be installed.
;
*****
;KGLWICMA=Tivoli Enterprise Monitoring Agent Framework
;KNTWICMA=Monitoring Agent for Windows OS
;KNT64CMA=Monitoring Agent for Windows OS (86-x64 only)
;KR2WICMA=Agentless Monitoring for Windows Operating Systems
;KR3WICMA=Agentless Monitoring for AIX Operating Systems
;KR4WICMA=Agentless Monitoring for Linux Operating Systems
;KR5WICMA=Agentless Monitoring for HP-UX Operating Systems
;KR6WICMA=Agentless Monitoring for Solaris Operating Systems
;KUMWICMA=Universal Agent
;KAC64CMA=32/64 Bit Agent Compatibility Package
;KUEWICMA=Tivoli Enterprise Services User Interface Extensions
```

5. Save the file and close the editor.

Take the following steps to run the silent uninstallation:

1. Open a DOS command prompt.
2. In the prompt, change to the directory containing this installation (where setup.exe and setup.ins files are located).
3. Run the setup as follows, specifying the parameters in the order listed:
`start /wait setup /z"/sfresponse_file" /s /f2"C:\temp\silent_setup.log"`

where:

- /z"/sf** Specifies the fully qualified name of the response file you edited. For example:
`/z"/sfC:\temp\myresponse.txt"`
- /s** Specifies that this is a silent installation. This causes nothing to be displayed during installation.
- /f2** Specifies the name of the InstallShield log file. If you do not specify this parameter, the default is to create Setup.log in the same location as the setup.iss file. In either case, the Setup program must be able to create and write to this file. If the specified directory does not exist, the Setup program cannot create and write to the file. Therefore, the specified directory path must exist.

When the uninstallation is complete, setup.exe returns to the command prompt.

If the uninstallation is unsuccessful, check the installation log in the *install_dir\InstallITM\product_nametime_stamp.log* directory. If all components and directories have been removed, the log will be in the root of the C: drive. The name of the file is the name of the product being uninstalled with date and time appended to the end of the product name.

Performing a silent uninstallation on a Linux or UNIX computer

Complete the following steps to uninstall IBM Tivoli Monitoring components and monitoring agents unattended (that is, without having to specify parameters interactively).

Procedure

1. Stop all agents and servers that you are going to remove. Use the `./cinfo -r` command to see a list of running process, and then use the `itmcmd` command to stop the appropriate processes.
2. Change to the installation directory bin directory:

```
cd install_dir/bin
```

3. To select individual components or agents, enter the following command:

```
uninstall.sh [-f] [-i] [-h install_directory] [product platformCode]
```

where

-f Forces delete, suppressing confirmation messages and prompts.

-i Ignores all running processes.

product

Is the two-letter code for the product to be uninstalled.

platformCode

Is the platform code for the product (such as aix513, sol286, hp11, and so forth: see Appendix D, “IBM Tivoli product, platform, and component codes,” on page 883).

Repeat the command for each agent or component you want to remove on the target computer.

4. To remove all components and agents enter the following command:

```
uninstall.sh REMOVE EVERYTHING
```

Results

When the uninstallation is complete, the `uninstall` command returns to the command prompt. Some messages may be written to the screen. There may be additional steps, depending on the component being uninstalled. For example, if you uninstall the Warehouse Proxy, the warehouse database is not removed and historical situations on the agent are not stopped (see “Uninstalling the Warehouse Proxy” on page 937).

If the uninstallation is unsuccessful, some messages may be written to the screen. See the installation log in the `install_dir/logs/product_nametime_stamp.log` directory or the for more information. If all components have been removed, the log is at the root.

What to do next

If for any reason, the UNIX uninstallation is not successful, run the following command to remove all Tivoli Monitoring directories:

```
rm -r install_dir
```

Uninstalling the event synchronization component

Use the following steps to uninstall the event synchronization from your event server.

Note: On Windows 2003 you must run the **change user /install** command from a command prompt before you begin. This puts the computer into the required “install” mode. After the uninstallation, run the **change user /execute** command to return the computer to its previous mode.

1. Set the Tivoli environment:

-  `C:\windows\system32\drivers\etc\Tivoli\setup_env.cmd`

or

```
C:\winnt\system32\drivers\etc\Tivoli\setup_env.cmd
```

- **Linux** **UNIX**

```
. /etc/Tivoli/setup_env.sh
```

2. Run the following uninstallation program:

- **Windows**

```
%BINDIR%\TME\TEC\OM_TEC\_uninst\uninstaller.exe
```

- **Linux** **UNIX**

```
$BINDIR/TME/TEC/OM_TEC/_uninst/uninstaller.bin
```

You can run this uninstallation program in silent mode (by running the program from the command-line with the **-silent** parameter) or in console mode (by using the **-console** parameter).

3. Follow the prompts in the uninstallation program.

When the uninstallation is completed, you can tell the installer what rule base should be loaded. If initial installation created a new rule base, the value shown in “Rule base name of rule base to be loaded on completion of this uninstall” will be Default, meaning that the Default rule base will be loaded. If the initial installation updated an existing rule base, that rule base name is provided as the value for “Rule base name of rule base to be loaded on completion of this uninstall”. You can override this value by typing in the name of the rule base you want to have loaded.

You can also tell the uninstaller to stop and restart the event server.

You can run the silent uninstallation using default processing or create a template to change the default values. The default processing will load the Default rule base (or the existing rule base was chosen during installation) and will *not* restart the TEC server.

To create and use a template:

1. Create the template:

- **Windows**

```
%BINDIR%\TME\TEC\OM_TEC\_uninst\uninstaller.exe -options-template  
itmeventsynchU.txt
```

- **Linux** **UNIX**

```
$BINDIR/TME/TEC/OM_TEC/_uninst/uninstaller.bin -options-template  
itmeventsynchU.txt
```

2. Modify the template as desired:

- To specify which rule base to load, modify the restartTECU.uRBN file.
- To automatically restart the event server, modify the restartTECU.restartTECU file.

3. Set the Tivoli environment.

4. Run the uninstallation program as follows:

- **Windows**

```
%BINDIR%\TME\TEC\OM_TEC\_uninst\uninstaller.exe -options  
itmeventsynchU.txt -silent
```

- **Linux** **UNIX**

```
$BINDIR/TME/TEC/OM_TEC/_uninst/uninstaller.bin -options  
itmeventsynchU.txt -silent
```

If your event server is running on an HP-UX computer, ensure that the `_uninst` and `_jvm` directories are successfully removed by the uninstallation program. If they are not, manually delete these directories.

Uninstalling event synchronization manually

Use the instructions for the appropriate operating system to remove event synchronization manually.

HP11

1. Stop the situation update forwarder long-running process if it is still running:
 - a. Find the long-running process using:

```
ps -ef
```
 - b. Use the kill command to remove the process:

```
kill -9 process_number
```
2. Run the following command to determine whether the operating system still knows that the event synchronization component is there:

```
swlist -v TecEvntSyncInstaller
```

If it is there but all the code has been deleted, or just the uninstaller is deleted, you can try this command:

```
swremove TecEvntSyncInstaller
```

3. If errors are returned saying that TecEvntSyncInstaller cannot be removed due to consistency or dependency checks, create a file named something like "remove_EvntSync.txt" and add these two lines:

```
enforce_dependencies=false  
enforce_scripts=false
```

Then run the **swremove** command as follows:

```
swremove -X remove_EvntSync.txt TecEvntSyncInstaller
```

The -X option tells the **swremove** command to ignore checks and dependencies and remove the files regardless.

4. Remove any event synchronization directories that are left behind.

Remove any directories found in OM_TEC including OM_TEC itself. OM_TEC is found in \$BINDIR/TME/TEC. To use \$BINDIR you must run the following command:

```
. /etc/Tivoli/setup_env.sh
```

If the installation was for Netcool/OMNIBus, remove files from the location indicated during installation.

Windows

1. Stop the situation update forwarder long running process if it is still running:
 - a. In the Control Panel, open **Administrative Tools**, then **Services**.
 - b. Find the Tivoli Situation Update Forwarder service, and right-click it and select **Stop**.
2. Go to operating system directory (C:\windows or C:\winnt) and open the vpd.properties file.
3. Remove all lines that have itmTecEvntSyncProduct, EvntSyncForwarder, itmTecEvntSyncLapComp or EvntSyncForwarderWin in them.
4. Remove any event synchronization directories that are left behind.

Remove any directories found in OM_TEC including OM_TEC itself. OM_TEC is found in %BINDIR%/TME/TEC. To use %BINDIR% you must run the C:\windows\system32\drivers\etc\Tivoli\setup_env.cmd command. If the installation was for Netcool/OMNIBus, remove the files from the location indicated during installation.

AIX

1. Stop the situation update forwarder long running process if it is still running:

- a. Find the long running process using:


```
ps -ef
```
- b. Use the kill command to remove the process:


```
kill -9 process_number
```
2. Go to operating system directory (this is typically /usr/lib/objrepos) and open the vpd.properties file.
3. Remove all lines that have itmTecEvtSyncProduct, EvntSyncForwarder, itmTecEvtSyncLapComp or EvntSyncForwarderWin in them.
4. Remove any event synchronization directories that remain.
Remove any directories found in OM_TEC including OM_TEC itself. OM_TEC is found in \$BINDIR/TME/TEC. To use \$BINDIR you must run the following command:


```
. /etc/Tivoli/setup_env.sh
```

If the installation was for Netcool/OMNIbus, remove the files from the location specified during installation.

Linux

1. Stop the situation update forwarder long running process if it is still running:
 - a. Find the long running process using:


```
ps -ef
```
 - b. Use the kill command to remove the process:


```
kill -9 process_number
```
2. Go to operating system directory (this is typically / or /root) and open the file vpd.properties.
3. Remove all lines that have itmTecEvtSyncProduct, EvntSyncForwarder, itmTecEvtSyncLapComp or EvntSyncForwarderWin in them.
4. Remove any event synchronization directories that are left behind.
Remove any directories found in OM_TEC including OM_TEC itself. OM_TEC is found in \$BINDIR/TME/TEC. To use \$BINDIR you must run the following command:


```
. /etc/Tivoli/setup_env.sh
```

If the installation was for OMNIbus, remove the files from the location specified during installation

Solaris

1. Stop the situation update forwarder long running process if it is still running:
 - a. Find the long running process using:


```
ps -ef
```
 - b. Use the kill command to remove the process:


```
kill -9 process_number
```
2. Run the following command to remove the situation update forwarder:


```
pkgrm -A ISitmTecE
```
3. Remove any event synchronization directories that are left behind.
Remove any directories found in OM_TEC including OM_TEC itself. OM_TEC is found in \$BINDIR/TME/TEC. To use \$BINDIR you must run the following command:


```
. /etc/Tivoli/setup_env.sh
```

If the installation was for OMNIbus, remove the files from the location specified during installation.

Uninstalling IBM Tivoli Monitoring components for a dashboard environment

The topics in this section provide information about uninstalling IBM Tivoli Monitoring components for a dashboard environment.

You can uninstall the IBM Infrastructure Management Dashboards for Servers, Tivoli Authorization Policy Server, and tivcmd Command Line Interface for Authorization Policy by using the IBM Installation Manager graphical user interface, console mode, or silent mode.

Uninstalling IBM Tivoli Monitoring components for a dashboard environment in console mode

Use the following steps to uninstall IBM Tivoli Monitoring components for a dashboard environment in console mode.

Before you begin

Note: The IBM WebSphere Application Server is restarted if IBM Infrastructure Management Dashboards for Servers is uninstalled.

Procedure

1. Start the IBM Installation Manager in console mode from the directory where it was installed. For example:

- **Windows** `C:\Program Files\IBM\Installation Manager\eclipse\tools\imcl.exe -c`
- **Linux** **UNIX** `/opt/IBM/InstallationManager/eclipse/tools/imcl -c`

2. The following menu is displayed:

```
Select:
  1. Install - Install software packages
  2. Update - Find and install updates and fixes to installed software packages
  3. Modify - Change installed software packages
  4. Roll Back - Revert to an earlier version of installed software packages
  5. Uninstall - Remove installed software packages

Other Options:
  L. View Logs
  S. View Installation History
  V. View Installed Packages
  -----
  P. Preferences
  -----
  E. Export Data for Problem Analysis
  A. About IBM Installation Manager
  -----
  X. Exit Installation Manager
```

Enter **5** *Uninstall - Remove installed software packages*

3. Select the package group that you want to uninstall.
 1. ☐ tivcmd Command Line Interface for Authorization Policy
 2. ☐ Core services in Jazz for Service Management

The package group **tivcmd Command Line Interface for Authorization Policy** contains the package for the tivcmd command line interface. The package group **Core services in Jazz for Service Management** contains the packages for IBM Infrastructure Management Dashboard for Servers and Tivoli Authorization Policy Server.

After you select a package number, IBM Installation Manager displays specific information about the package group and you select the packages to uninstall:

- a. If you selected the tivcmd Command Line Interface for Authorization Policy package group, select **tivcmd Command Line Interface for Authorization Policy** and enter N (Next) to continue.
- b. If you selected the Core services in Jazz for Service Management package group, select one or both packages to uninstall:
 - **Tivoli Authorization Policy Server 01.01.00.00**
 - **IBM Infrastructure Management Dashboards for Servers 1.1.0.0**
 Enter N (Next) to continue.
- c. Enter your WebSphere administrator-level user credentials (user name and password). The installer validates the entered credentials. Enter N (Next) to continue.
4. A summary screen displays the packages to be uninstalled. Select one of the following options:

Options:

G. Generate an Uninstall Response File
 B. Back, U. Uninstall, C. Cancel

- a. Enter G to record your choices in a silent response file.
- b. Enter U to proceed with the uninstallation.

Performing a silent uninstallation of IBM Tivoli Monitoring components for a dashboard environment

You can uninstall a package silently by using a response file.

Before you begin

You must create response file before you begin this procedure. For more information, see “Silent response file for uninstallation of IBM Tivoli Monitoring components for a dashboard environment” on page 946.

Note: The IBM WebSphere Application Server is restarted if IBM Infrastructure Management Dashboards for Servers is uninstalled.

Perform the following steps to uninstall a software package in silent mode:

Procedure

1. Open a command prompt and change to the tools sub directory in the directory where Installation Manager is installed. The default path for this directory is:
 - **Windows** C:\Program Files\IBM\InstallationManager\ eclipse\tools\
 - **Linux** **UNIX** /opt/IBM/InstallationManager/eclipse/tools/
2. Run the following command:
 - **Windows** imcl.exe input response_file -log log_file
 - **Linux** **UNIX** ./imcl input response_file -log log_file

Results

When the silent uninstall is complete, a status of 0 is returned. If the uninstall cannot be completed, a non-zero number is returned.

For more information about uninstalling a package silently by using a response file, go to http://pic.dhe.ibm.com/infocenter/install/v1r6/topic/com.ibm.silentinstall12.doc/topics/t_silent_response_file_uninstall.html.

Silent response file for uninstallation of IBM Tivoli Monitoring components for a dashboard environment

You must create response file before you can perform a silent uninstallation of IBM Tivoli Monitoring components for a dashboard environment.

The *IBM Tivoli Monitoring Dashboards for Servers and Authorization Policy Components DVD* or DVD image contains the following sample response file:

responsefiles\uninstall_ITMAp4JazzSM.xml

This file can be customized and used for silent uninstallation. For information about recording a new response file, see “Recording a response file” on page 858.

Sample response file:

```
<?xml version="1.0" encoding="UTF-8"?>
<!--
Sample uninstall response file.
This file has to be customized according to user's environment settings and then
in IBM Installation Manager Home directory execute following command:
IBMIM -s -input uninstall_ITMAp4JazzSM.xml -nosplash -acceptLicense -log <logfile>
-->

<agent-input>
<server>
</server>

<uninstall modify='false'>
<!-- IBM Infrastructure Management Dashboards for Servers.
Comment line below to skip component uninstallation.
-->
<offering id='com.ibm.tivoli.monitoring.osd' profile='Core services in Jazz for Service Management'
features='itmosd-install,itmosd-config'/>

<!-- IBM Tivoli Authorization Policy Server.
Comment line below to skip component uninstallation.
-->
<offering id='com.ibm.tivoli.kge.server.offering' profile='Core services in Jazz for Service Management'
features='server-install,server-config'/>

<!-- IBM Tivoli tivcmd Command Line Interface for Authorization Policy.
Comment line below to skip component uninstallation.
-->
<offering id='com.ibm.tivoli.kge.cli.offering' profile='tivcmd Command Line Interface for Authorization
Policy' features='kge.cli.feature'/>

</uninstall>

<profile id='Core services in Jazz for Service Management'>
<!-- Obligatory section when IBM Infrastructure Management Dashboards for Servers or
Authorization Policy Server is selected for uninstallation -->
<!-- DASH login credentials
Customize value below to provide DASH user name
-->
<data key='user.WAS_USER_NAME' value='tipadmin'/>
<!-- Encrypted DASH user password
Customize value below to provide DASH user password for your system. Value must be in encrypted form.
To encrypt the value, use IBM Installation Manager command line tool "imutilsc", stored IIM tools
subdirectory.
Example:
im.<platform>\tools>imutilsc encryptString <stringToEncryption>
Output of this command set as value in line below.
-->
<data key='user.WAS_PASSWORD' value='M6NQ+7Gb+GuoNQNg8N+Lkw=='/>
</profile>
</agent-input>
```

Uninstalling the Tivoli Authorization Policy Server or tivcmd Command Line Interface for Authorization Policy using the graphical user interface

Use the following steps to uninstall the Tivoli Authorization Policy Server or tivcmd Command Line Interface from your computer.

Procedure

1. Launch the IBM Installation Manger:
 - **Windows** Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager\eclipse.
 - **Linux** **UNIX** Execute the IBMIM binary under /opt/IBM/InstallationManger/eclipse.
2. On the Installation Manager window, click **Uninstall**.
3. Depending on what package you are uninstalling, select either the **Tivoli Authorization Policy Server** OR the **tivcmd Command Line Interface for Authorization Policy** check box and click **Next**.
4. If you are uninstalling the Tivoli Authorization Policy Server, enter the WebSphere Application Server **user name** and **password** in the fields provided. The login credentials must be validated before you click **Next**.
5. Review the uninstall summary information and click **Uninstall**.
The uninstall completion panel is displayed.

Results

What to do next

If you are uninstalling both the Installation and Configuration features, perform the following checks:

- Check that there are no errors in the deployPolicyServer.log file that is located in <install_dir>\JazzSM\logs.
- Check that the AuthPolicyServer folder is deleted in <install_dir>\JazzSM.
- Check that the PolicyServer.war folder is deleted in <install_dir>\JazzSM\profile\installedApps\localhostNode01Cell\isc.ear\PolicyServer.war.

If you are uninstalling the Installation feature only, perform the following check:

- Check that the AuthPolicyServer folder is deleted in <install_dir>\JazzSM.

Uninstalling the IBM Infrastructure Management Dashboards for Servers using the graphical user interface

Use the following steps to uninstall the IBM Infrastructure Management Dashboards for Servers from your computer.

Before you begin

Note: The installer will restart the IBM WebSphere Application Server when you uninstall IBM Infrastructure Management Dashboards for Servers.

Procedure

1. Launch the IBM Installation Manger:
 - **Windows** Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where IBM Installation Manager was installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager.

- **Linux** **UNIX** Execute the IBMIM binary under /opt/IBM/InstallationManger/eclipse.
2. Click **Uninstall** to see the installed packages on the Dashboard Application Services Hub server, and select the **Infrastructure Management Dashboards for Servers** check box. The features that are associated with the infrastructure management dashboards are also selected by default. After you click **Next**, the next panel displays with the Common Configurations and fields for entering your WebSphere Application Server credentials.
 3. Log on to the local Dashboard Application Services Hub server by entering your administrator-level user credentials, such as smadmin and its password, in the **User name** and **Password** fields, and click **Validate**. After your credentials are validated and you click **Next**, the summary of packages that you have chosen to uninstall is displayed.
 4. Review the summary of your selections and click **Uninstall** to begin the uninstallation process. This is your opportunity to click **Back** and make any changes before proceeding. After uninstallation is complete, the packages are displayed with a message confirming that they have been uninstalled. You can now review the log files.
 5. Click **Finish** to finalize the uninstallation.

Appendix I. Additional resources

The following sections include resources and information you can use during your Tivoli Monitoring planning and deployment.

IBM Tivoli Monitoring 6 Welcome Kit

The purpose of this guide is to provide guidelines and reference materials to assist you in getting yourself familiar with the Tivoli Monitoring version 6 product.

It was produced with the following objectives in mind:

- Help you to effectively use Tivoli Monitoring
- Obtaining Tivoli Monitoring assistance online
- Provide additional resources for working with Tivoli Monitoring
- Introduce you to the Tivoli Monitoring Support
- Provide an image containing documents and presentations on Tivoli Monitoring

Documents can be viewed at <http://www.ibm.com/support/docview.wss?rs=2366&uid=swg21253835>.

General education and support Web sites

IBM Tivoli Monitoring Documentation Central

IBM Tivoli Monitoring Documentation Central (<http://www.ibm.com/tivoli/documentation>)

ISM Library

IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary>)

IBM Service Management Connect

IBM Service Management Connect (<http://www.ibm.com/developerworks/servicemanagement>)

IBM Service Management Connect > Application Performance Management (<http://www.ibm.com/developerworks/servicemanagement/apm>)

IBM Service Management Connect > Application Performance Management > Tivoli Monitoring Forum (<http://www.ibm.com/developerworks/forums/rss/rssmessages.jspa?forumID=796>)

IBM Service Management Connect > Jazz for Service Management (<https://www.ibm.com/developerworks/servicemanagement/iosdp/index.html>)

Tivoli Common Reporting

IBM Tivoli Common Reporting community (<http://www.ibm.com/developerworks/spaces/tcr>)

Main Tivoli Page

<http://www.ibm.com/tivoli>

This is the main Tivoli page (formerly www.tivoli.com)

Main Tivoli Support Page

<http://www.ibm.com/software/sysmgmt/products/support/>

This is the main Tivoli Support page. All Tivoli products are listed in the drop down box in the middle. Choosing a product takes you to the support page dedicated to the product you need help with.

Main Tivoli Education Page

<http://www.ibm.com/software/tivoli/education/>

IBM Tivoli offers a variety of course types such as: online, on-site and instructor-led education. The courses cover aspects of the Tivoli software portfolio.

Main IBM Redbook Page

<http://www.redbooks.ibm.com/redbooks.nsf/redbooks/>

IBM Redbooks are developed and published by IBM's International Technical Support Organization. They deliver skills, technical know-how, and materials to technical professionals of IBM, Business Partners, and users and to the marketplace generally.

IBM Certified Advanced Deployment Professional - IBM Service Management Service Availability and Performance Management

<http://www-03.ibm.com/certify/certs/24109303.shtml>

An IBM Certified Advanced Deployment Professional - IBM Service Management Service Availability and Performance Management is an individual who has demonstrated a higher level of implementation knowledge and skill both in breadth and in depth in the IBM Tivoli Service Availability and Performance Management solutions area.

Product documentation and IBM Redbooks

IBM Tivoli Monitoring product documentation

<http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/>

Infrastructure Solutions: Building a Smart Bank Operating Environment, SG24-7113

<http://www.redbooks.ibm.com/abstracts/sg247113.html?Open>

Implementing OMEGAMON XE for Messaging 6.0, SG24-7357

<http://www.redbooks.ibm.com/abstracts/sg247357.html?Open>

Best Practices for SOA Management, REDP-4233

<http://www.redbooks.ibm.com/abstracts/redp4233.html?Open>

Deployment Guide Series: IBM Tivoli Monitoring 6.2, SG24-7444

<http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg247444.html?Open>

Getting Started with IBM Tivoli Monitoring 6.1 on Distributed Environments, SG24-7143

<http://www.redbooks.ibm.com/abstracts/sg247143.html?Open>

Tivoli Management Services Warehouse and Reporting, SG24-7290

<http://www.redbooks.ibm.com/abstracts/sg247290.html?Open>

IBM Tivoli OMEGAMON XE V3.1 Deep Dive on z/OS, SG24-7155

<http://www.redbooks.ibm.com/Redbooks.nsf/ad6437adb3f17484852568dd006f956e/b8e023248b2d90718525706c00609acd?OpenDocument>

IBM Tivoli Monitoring: Implementation and Performance Optimization for Large Scale Environments, SG24-7443

<http://www.redbooks.ibm.com/redpieces/abstracts/sg247443.html?Open>

Education offerings

A listing of all the current Tivoli Monitoring training can be found at http://www-304.ibm.com/jct03001c/services/learning/ites.wss/zz/en?pageType=tp_search.

Support Technical Exchange (STE) Seminars

Expand your technical understanding of your current Tivoli products, in a convenient format hosted by IBM Tivoli Worldwide Support and Services. These live seminars are support oriented discussions of product information, deployment and trouble-shooting tips, common issues, problem solving resources and other support and service recommendations. Tivoli engineers and consultants who are subject matter experts for the product(s) discussed lead each STE. Each STE is recorded and playback is available at

anytime. To attend a live STE or review a previously recorded STE go to http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html.

Service offerings

There are several Services offerings for the Tivoli Monitoring product. Access the Services offerings and additional details on some of the offerings at the following link:

<http://www-01.ibm.com/software/tivoli/services/consulting/offerings.html>.

Other resources

AA&BSM Enablement Best Practices website

<http://www.ibm.com/software/tivoli/features/monitoring-best-practices/index.html>

Tivoli AA&BSM Technical Exchange Wiki

IBM Tivoli AA&BSM Technical Exchange (<http://www.ibm.com/developerworks/wikis/display/aabsmenbl/Home>)

IBM Tivoli Monitoring 6 Forum

http://www.ibm.com/developerworks/forums/dw_forum.jsp?forum=796&cat=15

Appendix J. Documentation library

This appendix contains information about the publications related to IBM Tivoli Monitoring and to the commonly shared components of Tivoli Management Services.

These publications are listed in the following categories:

- IBM Tivoli Monitoring library
- Related publications

For information about accessing and using the publications, select **Using the publications** in the **Contents** pane of the IBM Tivoli Monitoring and OMEGAMON XE Information Center at <http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/index.jsp>.

To find a list of new and changed publications, click **What's new** on the Welcome page of the IBM Tivoli Monitoring and OMEGAMON XE Information Center. To find publications from the previous version of a product, click **Previous versions** under the name of the product in the **Contents** pane.

IBM Tivoli Monitoring library

The following publications provide information about IBM Tivoli Monitoring and about the commonly shared components of Tivoli Management Services:

- *Quick Start Guide*
Introduces the components of IBM Tivoli Monitoring.
- *Installation and Setup Guide*, SC22-5445
Provides instructions for installing and configuring IBM Tivoli Monitoring components on Windows, Linux, and UNIX systems.
- *Program Directory for IBM Tivoli Management Services on z/OS*, GI11-4105
Gives instructions for the SMP/E installation of the Tivoli Management Services components on z/OS.
- *High Availability Guide for Distributed Systems*, SC22-5455
Gives instructions for several methods of ensuring the availability of the IBM Tivoli Monitoring components.
- *IBM Tivoli zEnterprise Monitoring Agent Installation and Configuration Guide*, SC14-7358
Provides instructions for installing and configuring Tivoli zEnterprise monitoring agent components on Windows, Linux, and UNIX systems. Also includes migration and backup information, Enterprise Common Collector troubleshooting, Hardware Management Console configuration, and use of the command line interface or APIs to customize the collector. This guide complements the *Tivoli zEnterprise Monitoring Agent User's Guide*.
- *Administrator's Guide*, SC22-5446
Describes the support tasks and functions required for the Tivoli Enterprise Portal Server and clients, including Tivoli Enterprise Portal user administration.
- *Command Reference*, SC22-5448
Provides detailed syntax and parameter information, as well as examples, for the commands you can use in IBM Tivoli Monitoring.
- *Messages*, SC22-5450
Lists and explains messages generated by all IBM Tivoli Monitoring components and by z/OS-based Tivoli Management Services components (such as Tivoli Enterprise Monitoring Server on z/OS and TMS:Engine).
- *Troubleshooting Guide*, GC22-5449

Provides information to help you troubleshoot problems with the software.

- Tivoli Enterprise Portal online help

Provides context-sensitive reference information about all features and customization options of the Tivoli Enterprise Portal. Also gives instructions for using and administering the Tivoli Enterprise Portal.

- *Tivoli Enterprise Portal User's Guide*, SC22-5447

Complements the Tivoli Enterprise Portal online help. The guide provides hands-on lessons and detailed instructions for all Tivoli Enterprise Portal features.

- *Agent Builder User's Guide*, SC32-1921

Explains how to use the Agent Builder for creating monitoring agents and their installation packages, and for adding functions to existing agents.

- *Performance Analyzer User's Guide*, SC27-4004

Explains how to use the Performance Analyzer to understand resource consumption trends, identify problems, resolve problems more quickly, and predict and avoid future problems.

- *IBM Tivoli zEnterprise Monitoring Agent User's Guide*, SC14-7359

Complements the Tivoli zEnterprise monitoring agent online help. The guide provides reference information about the interface, usage scenarios, agent troubleshooting information, and information about Tivoli Common Reporting reports. This guide complements the *Tivoli zEnterprise Monitoring Agent Installation and Configuration Guide*.

Documentation for the base agents

If you purchased IBM Tivoli Monitoring as a product, you received a set of base monitoring agents as part of the product. If you purchased a monitoring agent product (for example, an OMEGAMON XE product) that includes the commonly shared components of Tivoli Management Services, you did not receive the base agents.

The following publications provide information about using the base agents.

- Operating system agents:
 - *Windows OS Agent User's Guide*, SC22-5451
 - *UNIX OS Agent User's Guide*, SC22-5452
 - *Linux OS Agent User's Guide*, SC22-5453
 - *IBM i Agent User's Guide*, SC22-5454
- Agentless operating system monitors:
 - *Agentless Monitoring for Windows Operating Systems User's Guide*, SC23-9765
 - *Agentless Monitoring for AIX Operating Systems User's Guide*, SC23-9761
 - *Agentless Monitoring for HP-UX Operating Systems User's Guide*, SC23-9763
 - *Agentless Monitoring for Solaris Operating Systems User's Guide*, SC23-9764
 - *Agentless Monitoring for Linux Operating Systems User's Guide*, SC23-9762
- Warehouse agents:
 - *Warehouse Summarization and Pruning Agent User's Guide*, SC22-5457
 - *Warehouse Proxy Agent User's Guide*, SC22-5456
- System P agents:
 - *AIX Premium Agent User's Guide*, SA23-2237
 - *CEC Base Agent User's Guide*, SC23-5239
 - *HMC Base Agent User's Guide*, SA23-2239
 - *VIOS Premium Agent User's Guide*, SA23-2238
- Other base agents:
 - *Tivoli Log File Agent User's Guide*, SC14-7484

Related publications

For information about related products and publications select **OMEGAMON XE shared publications** or other entries in the **Contents** pane of the IBM Tivoli Monitoring and OMEGAMON XE Information Center at <http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/index.jsp>.

Other sources of documentation

You can also obtain technical documentation about IBM Tivoli Monitoring and related products from the following sources:

- Service Management Connect (SMC)

For introductory information about SMC, see IBM Service Management Connect (<http://www.ibm.com/developerworks/servicemanagement>).

For information about Tivoli products, see the Application Performance Management community on SMC at IBM Service Management Connect > Application Performance Management (<http://www.ibm.com/developerworks/servicemanagement/apm>).

Connect, learn, and share with Service Management professionals. Get access to developers and product support technical experts who provide their perspectives and expertise. Using SMC, you can:

- Become involved with transparent development, an ongoing, open engagement between external users and developers of Tivoli products where you can access early designs, sprint demos, product roadmaps, and pre-release code.
- Connect one-on-one with the experts to collaborate and network about Tivoli and Integrated Service Management.
- Benefit from the expertise and experience of others using blogs.
- Collaborate with the broader user community using wikis and forums.

- Tivoli wikis

IBM Service Management Connect > Application Performance Management (<http://www.ibm.com/developerworks/servicemanagement/apm>) includes a list of relevant Tivoli wikis that offer best practices and scenarios for using Tivoli products, white papers contributed by IBM employees, and content created by customers and business partners.

Two of these wikis are of particular relevance to IBM Tivoli Monitoring:

- The IBM Tivoli Monitoring Wiki (<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/Tivoli%20Monitoring/page/Home>) provides information about IBM Tivoli Monitoring and related distributed products, including IBM Tivoli Composite Application Management products.
- The Tivoli System z Monitoring and Application Management Wiki provides information about the OMEGAMON XE products, NetView for z/OS, Tivoli Monitoring Agent for z/TPF, and other System z monitoring and application management products.

- IBM Integrated Service Management Library

<http://www.ibm.com/software/brandcatalog/ismlibrary/>

IBM Integrated Service Management Library is an online catalog that contains integration documentation and other downloadable product extensions.

- Redbooks

<http://www.redbooks.ibm.com/>

IBM Redbooks and Redpapers include information about products from platform and solution perspectives.

- Technotes

Technotes provide the latest information about known product limitations and workarounds. You can find Technotes through the IBM Software Support Web site at <http://www.ibm.com/software/support/>.

Appendix K. Support for problem solving

This section describes the options for obtaining support for IBM software products.

If you have a problem with your IBM software, you want to resolve it quickly. You can obtain support for IBM software products through the following options:

- “Using IBM Support Assistant”
- “Obtaining fixes” on page 958
- “Receiving weekly support updates” on page 958
- “Contacting IBM Software Support” on page 959

Using IBM Support Assistant

The IBM Support Assistant is a free, stand-alone application that you can install on most workstations and also use to perform remote troubleshooting of other workstations.

The IBM Support Assistant saves you the time it takes to search product, support, and educational resources. You can enhance the application by installing product-specific add-ons for the IBM products you use. Several troubleshooting features are provided, including the ability to perform guided troubleshooting to aid in problem resolution and the ability to collect diagnostic information. The collected diagnostic information can then be used to self-diagnose the problem, or it can be included in an *Electronic Service Request* (ESR) submitted to IBM Support engineers. The ESR tool is used to open, update, and report on PMRs (Problem Management Records) online. See <http://www.ibm.com/software/support/help.html> for assistance in using the ESR tool.

For more information, and to download the IBM Support Assistant, see <http://www.ibm.com/software/support/isa>. Currently, the add-on for this product is supported by IBM Support Assistant V4.0.1, or later.

Installing the IBM Support Assistant add-on

Install the IBM Support Assistant add-on for the IBMTivoli Monitoring product that you are using.

About this task

After you download and install the IBM Support Assistant, follow these steps to install the IBM Support Assistant add-on for the IBMTivoli Monitoring product that you are using:

Procedure

1. Start the IBM Support Assistant application.
2. From the **File** → **Preferences** → **Updater preferences** menu, provide the URL to update the site under **Specify an Update Site** → **Location**.
3. Select **http** from the list.
4. Validate the site and click **OK** to confirm changes.
5. Run **Update** → **Find new** → **Product Add-ons**.
6. Select the appropriate plug-in
7. Read the license and description, and if you comply, select **I accept the terms in the license agreements** and click **Next**.
8. Click **Finish** to proceed with the installation, and when prompted, restart the IBM Support Assistant to complete the installation.

Collection of diagnostic files for self-diagnosis only

You can collect diagnostic files and include them in an ESR that can be sent to IBM Support engineers.

About this task

To perform the collection of diagnostic files for self-diagnosis only, complete the following steps:

Procedure

1. Start the IBM Support Assistant application.
2. From the Home screen, select **Analyze Problem**.
3. In the Select A Collector dialog box, expand the appropriate product name, and select the agent for which you want to collect diagnostic information. Choose **Add**.
4. After the agent or agents are added to the Collector Queue, choose **Collect All** to begin the collection.
5. Enter the information requested in the dialog boxes.
6. The final dialog box requests whether or not you want to upload the collection file to IBM Support or another FTP location. If you only want to view the collected files on your computer, choose **Do Not FTP the Logs**.
7. The collection has finished. You can view the collected files by clicking the compressed file in the **Collector Status** dialog box.

Obtaining fixes

Find a product fix to resolve your problem.

About this task

A product fix might be available to resolve your problem. To determine which fixes are available for your Tivoli software product, follow these steps:

Procedure

1. Go to the IBM Software Support Web site at <http://www.ibm.com/software/support>.
2. Under **Select a brand and/or product**, select **Tivoli**.
3. Click the right arrow to view the Tivoli support page.
4. Use the **Select a category** field to select the product.
5. Select your product and click the right arrow that shows the **Go** hover text.
6. Under **Download**, click the name of a fix to read its description and, optionally, to download it. If there is no **Download** heading for your product, supply a search term, error code, or APAR number in the field provided under **Search Support (this product)**, and click the right arrow that shows the **Go** hover text.

Results

For more information about the types of fixes that are available, see the *IBM Software Support Handbook* at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

Receiving weekly support updates

You can receive weekly e-mail notifications about fixes and other software support news.

About this task

To receive weekly e-mail notifications about fixes and other software support news, follow these steps:

Procedure

1. Go to the IBM Software Support Web site at <http://www.ibm.com/software/support>.
2. Click **My support** in the far upper-right corner of the page under **Personalized support**.
3. If you have already registered for **My support**, sign in and skip to the next step. If you have not registered, click **register now**. Complete the registration form using your e-mail address as your IBM ID and click **Submit**.
4. The **Edit profile** tab is displayed.
5. In the first list under **Products**, select **Software**. In the second list, select a product category (for example, **Systems and Asset Management**). In the third list, select a product sub-category (for example, **Application Performance & Availability** or **Systems Performance**). A list of applicable products is displayed.
6. Select the products for which you want to receive updates.
7. Click **Add products**.
8. After selecting all products that are of interest to you, click **Subscribe to email** on the **Edit profile** tab.
9. In the **Documents** list, select **Software**.
10. Select **Please send these documents by weekly email**.
11. Update your e-mail address as needed.
12. Select the types of documents you want to receive.
13. Click **Update**.

What to do next

If you experience problems with the **My support** feature, you can obtain help in one of the following ways:

Online

Send an e-mail message to erchelp@ca.ibm.com, describing your problem.

By phone

Call 1-800-IBM-4You (1-800-426-4968).

Contacting IBM Software Support

IBM Software Support provides assistance with product defects. The easiest way to obtain that assistance is to open a PMR or ETR directly from the IBM Support Assistant.

Before contacting IBM Software Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenance contract that you need depends on the type of product you have:

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus®, and Rational® products, as well as DB2 and WebSphere products that run on Windows or UNIX operating systems), enroll in Passport Advantage in one of the following ways:

Online

Go to the Passport Advantage website at http://www-306.ibm.com/software/howtobuy/passportadvantage/pao_customers.htm.

By phone

For the phone number to call in your country, go to the IBM Software Support website at <http://techsupport.services.ibm.com/guides/contacts.html> and click the name of your geographic region.

- For customers with Subscription and Support (S & S) contracts, go to the Software Service Request website at <https://techsupport.services.ibm.com/ssr/login>.

- For customers with Linux, iSeries, pSeries, zSeries, and other support agreements, go to the IBM Support Line website at <http://www.ibm.com/services/us/index.wss/so/its/a1000030/dt006>.
- For IBM eServer™ software products (including, but not limited to, DB2 and WebSphere products that run in zSeries, pSeries, and iSeries environments), you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for eServer software products, go to the IBM Technical Support Advantage website at <http://www.ibm.com/servers/eserver/techsupport.html>.

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. From other countries, go to the contacts page of the *IBM Software Support Handbook* on the web at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click the name of your geographic region for phone numbers of people who provide support for your location.

Determining the business impact of a problem

Assess the severity and business impact of the problem that you are reporting.

When you report a problem to IBM, you are asked to supply a severity level. Use the following criteria to understand and assess the business impact of the problem that you are reporting:

Severity 1

The problem has a *critical* business impact. You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.

Severity 2

The problem has a *significant* business impact. The program is usable, but it is severely limited.

Severity 3

The problem has *some* business impact. The program is usable, but less significant features (not critical to operations) are unavailable.

Severity 4

The problem has *minimal* business impact. The problem causes little impact on operations, or a reasonable circumvention to the problem was implemented.

Describing problems and gathering information

When describing a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- Which software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can you re-create the problem? If so, what steps were performed to re-create the problem?
- Did you make any changes to the system? For example, did you make changes to the hardware, operating system, networking software, and so on.
- Are you currently using a workaround for the problem? If so, be prepared to explain the workaround when you report the problem.

Submitting problems

Submit problems to IBM Software Support

You can submit your problem to IBM Software Support in one of two ways:

Online

Click **Submit and track problems** on the IBM Software Support site at <http://www.ibm.com/software/support/probsub.html>. Type your information into the appropriate problem submission form.

By phone

For the phone number to call in your country, go to the contacts page of the *IBM Software Support Handbook* at <http://techsupport.services.ibm.com/guides/contacts.html> and click the name of your geographic region.

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the Software Support Web site daily, so that other users who experience the same problem can benefit from the same resolution.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law :

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2013. All rights reserved.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of The Minister for the Cabinet Office, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's user name for purposes of session management, authentication, and single sign-on configuration. These cookies cannot be disabled.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Index

Special characters

/etc/hosts file 805
<bind> element 870
<connection> element 870
<gateway> element 870
<interface> element 870
<portpool> element 870
<zone> element 870

A

AC (agent compatibility) component 265
 errors 262, 265
 monitoring server
 configuration file 262
 environment file 262
 KBBENV configuration file 262
access plan 494
activating a firewall gateway 869
Active Directory, Microsoft 137
adding application support
 for nonbase agents 279
 itmcmd support command 209
 Linux desktop client 291
 Linux monitoring server 161, 208, 284
 Linux portal server 288
 to a remote monitoring server from
 Linux or UNIX 294
 to a remote monitoring server from
 Windows 292
 UNIX monitoring server 161, 208, 284
 Windows desktop client 290
 Windows monitoring server 161, 281
 Windows portal server 287
advanced configuration
 Linux 439
 UNIX 439
agent autonomy
 background information 55
 event forwarding 55
Agent Builder Agents
 deployment task estimates 71
Agent Compatibility Package 265
agent deployment 371
 deploying OS agents 375
 managing agent depot 374
 OS agents, deploying 375
 populating agent depot 371
 sharing an agent depot 375
 tacmd createNode command 375
agent depot 371
 DEPOTHOME environment
 variable 374
 location 374
 managing 374
 populating through install 372, 373
 populating with tacmd addBundles
 command 373
agent depot (*continued*)
 sharing 375
agentless monitoring 56
 AIX 56
 deployment options 60
 documentation for 61
 HP-UX 56
 Linux 56
 platforms monitored 56
 problem-diagnosis tools for 62
 Solaris 56
 Windows 56
AGENTPRI parameter 497
agents
 application agents 11
 configuring 417
 deploying 371
 deploying OS agents 375
 IBM Tivoli Monitoring for
 Applications: mySAP Agent 156
 IBM Tivoli Monitoring for Databases:
 DB2 for Linux, UNIX, and Windows
 Agent 156
 IBM Tivoli Monitoring for Databases:
 Oracle Agent 156
 IBM Tivoli Monitoring for Databases:
 Sybase Server Agent 156
 IBM Tivoli Monitoring: UNIX OS
 Agent 156
 IBM Tivoli Monitoring: Windows OS
 Agent 156
 non-OS agents 11
 operating system (OS) agents 11
 OS agents 11
 product codes 883
 removing through the portal 937
 self-describing 204
 starting 420
 stopping 420
 uninstalling OMEGAMON 935
 uninstalling on Linux 935
 uninstalling on UNIX 935
 uninstalling on Windows 934
aggregate data, amount of 511
AIX
 configuring the portal server 243, 250
 installing the portal server 242
 portal server installation 241
application support
 adding to a remote monitoring server
 from Linux or UNIX 294
 adding to a remote monitoring server
 from Windows 292
 for agents on Base DVDs 278
 for an IBM Tivoli Composite
 Application Manager agent 275
 for nonbase agents 279
 application support media 277
application support, enabling
 enabling 275
application support, installing
 catalog and attribute files 275
 seed data 275
 SQL files 275
archdsc.tbl 137
architecture codes
 Linux 137
 UNIX 137
ASLHEAPSZ parameter 497
asterisk (*) wildcard character 501
asymmetric cryptography 454
asynchronous page cleaners 498
attribute group
 aggregate data, amount of 511
 detailed data, amount of 511
 detailed records per day, number
 of 510
 hard disk drive footprint 511
attributes
 event synchronization event
 class 818
authenticating passwords 139
authenticating users 139
authentication
 using external LDAP registry 139
 using local registry 139
Authorization Policy Server
 modifying 328
 rollback 336
 updating 334
 verifying 329
authorizing users 139
Automation Server
 configuring on Windows 414
 installing on Windows 220
autonomous OS agents 399
autostart scripts
 changes to 130

B

backing up a current installation
 UNIX or Linux 154
 Windows 152
backing up the portal server
 database 155
backing up the Tivoli Data Warehouse
 database 156
barrier firewall 861
base monitoring services 5
benchmarking tool 501, 502
Big Local Zone 135
Brio 13
browser
 configuring location on UNIX and
 Linux 364
 configuring location on
 Windows 363
browser client
 starting 366, 367

- buffer pools
 - definition and purpose 492
 - memory allocation 492
- BufferFlushRate parameter 764, 821
- buffering
 - advantages 492
 - double 491
- buildpresentation.bat 159
- bulk agent deployment 382

C

- cache, clearing 494
- caching
 - dynamic SQL 494
 - Windows system 491
- cardinality 489
- cat and atr files
 - location of on Windows 292
- catalog and attribute (cat and atr) files 275
- change user /execute command 130
- changing configuration settings 100
- CHNGPGS_THRESH parameter 498
- cinfo command 850, 886
- circular logging, definition and purpose 492
- Citrix client 130
- class, Sentry2_0_Base 126, 804
- Classic REORG 493
- clustered index 489
- coexistence with version 6.3 158
- Cold Backup 44
- COLLECT STATISTICS clause 494
- command
 - tacmd createNode 377
- commands 202
 - cgabge user/execute 130
 - itmcmd agent 271, 420
 - itmcmd config 851
 - itmcmd manage 412
 - itmcmd server 209, 420
 - itmcmd support 209
 - sitconfig.sh 822
 - sitconfuser.sh 822
 - tacmd addBundles 373
 - tacmd createNode 375
- Common Installer 843
- common services components 5
- communications
 - securing 138
 - security options for 137
- components
 - Agents DVD 11, 60, 277
 - CD 11
 - deciding which to upgrade 149
 - event synchronization 13
 - Infrastructure DVD 11, 277
 - monitoring agents 11
 - monitoring server 10
 - Performance Analyzer Warehouse Agent 11
 - portal client 11
 - portal server 11
 - Summarization and Pruning Agent 13
 - Tivoli Data Warehouse 13
- components (*continued*)
 - Tivoli Enterprise Portal 11
 - Tivoli Enterprise Portal Server 11
 - Tivoli Performance Analyzer 18, 33
 - Warehouse Proxy 13
- configuration 101
- configuration fields
 - Netcool/OMNIBus event synchronization 740
- configuration variables, agent 887
- configuring
 - agents 417
 - Automation Server 414
 - hub monitoring server, Linux 202
 - hub monitoring server, UNIX 202
 - IPv6 communication 441
 - Linux, silent 851
 - monitoring agents 417
 - monitoring agents on Linux 269
 - monitoring agents on UNIX 269
 - monitoring server 412
 - monitoring server, Linux 216
 - password authentication 139
 - port number assignments 421
 - portal server 453
 - portal server on Linux or AIX 243, 250
 - remote monitoring server, UNIX 216
 - response file, using 851
 - Tivoli Enterprise Monitoring Automation Server, Linux 224, 415
 - Tivoli Enterprise Monitoring Automation Server, UNIX 224, 415
 - UNIX, silent 851
- Configuring
 - browsers 342
 - clients 342
 - JREs 342
- configuring a firewall gateway 869
- configuring Netcool/OMNIBus
 - ObjectServer
 - error event flow 773
 - for program execution from scripts 763
 - updating the Netcool/OMNIBus database schema 753
- configuring the desktop client 343
- connection
 - inbound 495
 - outbound 495
 - pooling 495
- containers, definition and purpose 490
- control block information 498
- controllers, using separate 495
- cookies 965
- copyright 963
- CPU utilization 495
- creating a workspace 665
- Crystal Reports 13
- CT/Service Console 422
- CTIRA_HEARTBEAT environment variable 426
- CURRENT DEGREE special register 498
- customer support 959

D

- daily health checks 122
- dashboard
 - verifying the dashboard data provider is enabled 428
- data integrity, importance of logs 492
- data size, calculating 495
- data warehouse
 - creating database objects for 525
 - creating functions for 525
 - creating indexes for 525
 - creating table inserts for 525
 - creating tables for 525
 - creating views for 525
- data warehouse ratio 495
- Database Managed Space (DMS) table space 490
- database objects, creating manually 525
- database parameters 498
- database size, estimating 512
- databases
 - configuration parameters 492
 - heap 492
 - reorganization 493
- databases, supported
 - Tivoli Data Warehouse 82
- DB2 Control Center 497
- DB2 Database for Linux, UNIX, and Windows
 - upgrading and migrating 147
- DB2 for Linux, UNIX and Windows data warehouse
 - ODBC connection 551
- DB2 for Linux, UNIX, and Windows
 - creating a table space 514
 - IBMDEFAULTGROUP 514
 - increasing the size 514
- DB2 for Linux, UNIX, and Windows Enterprise Edition 82
- DB2BATCH 502
- db2batch tool 501
- db2empfa command 491
- DB2NTNOCACHE option 491
- DB2SET command 501
- DBHEAP parameter 498
- default certificate 142
- default event destination, defining 764
- Default Small Local Zone 135
- defining a portal server interface on Linux or UNIX 877
- defining a portal server interface on Windows 876
- DEGREE bind option 498
- Deploy Status attribute group 385
- Deploy Summary attribute group 385
- deploying
 - non-OS agents 377
 - through the command-line 378
 - through the Tivoli Enterprise Portal 378
- deployment 97
 - configuring your warehouse agents 101
 - installing additional agents 102
 - installing infrastructure components 97
 - installing your first agents 101

- deployment (*continued*)
 - planning 103
 - post-installation checklist 101
 - pre-installation checklist 97
- Deployment Status workspaces 384
- deployment task estimates 66
 - Agent Builder Agents 71
 - agents 69
 - deployment 71
 - fix packs 71
 - policies and workflows 70
 - situations 70
 - skills transfer 71
 - Summarization and Pruning Agent 68
 - Tivoli Enterprise Console integration 68
 - Warehouse Proxy Agent 68
 - Windows and UNIX environments 67
 - workspaces 70
 - z/OS servers 67
- DEPOTHOME environment variable 374
- DEPOTHOME keyword
 - KBBENV configuration file 36
- Derby embedded portal server database 82
- desktop client
 - adding application support on Linux 291
 - adding application support on Windows 290
 - configuring 343
 - configuring on Linux 275
 - creating a shortcut to launch using Web Start 369
 - downloading from portal server 368
 - installing 272
 - installing on Linux from installation media 273
 - installing on Windows from installation media 272
 - launching from IBM Java Control Panel 368
 - logs, location of 367
 - planning worksheet, Linux 839
 - planning worksheet, Windows 838
 - starting 366
 - using Web Start from the command-line 368
- detailed data, amount of 511
- detailed records per day, number of 510
- developerWorks 955
- DFT_DEGREE parameter 498
- diagnostic files 958
- digital certificates 454
- disk I/O, minimizing 492
- disk requirements for Tivoli Data Warehouse 513
- disk space, calculating 495
- disk storage requirements 85
- disks, drives 495
- downloading the desktop client using Java Web Start 368
- DSS ratio 495
- dynamic affinity, agent 160

E

- education 949
 - offerings 950
- EGG1 encryption scheme 142
- EIB 24
- EIF Probe 31
- EIF Probe Rules
 - updating 772
- EIF probe, configuring 760
- Enable Multipage File Allocation tool 491
- enabling
 - self-describing agent capability 204, 397
- enabling application support for nonbase agents 279
- enabling event forwarding on the monitoring server 764
- enabling tracing for the JRE 367
- Enterprise Information Base 24
- environment
 - changing configuration settings 100
 - configuration checklist 98
 - customizing your environment 99
 - installing components 97
 - maintenance 121
- environment variables
 - CTIRA_HEARTBEAT 426
 - KDC_FAMILIES 423, 425
 - KDE_TRANSPORT 423, 425
- ephemeral keyword 863
- ephemeral pipe
 - implementing 863
- estimating warehouse database size 510
 - aggregate data, amount of 511
 - amount of detailed data 511
 - detailed records per day, number of 510
 - determining the size 512
 - hard disk drive footprint 511
 - worksheet 512
- event integration scenarios with Tivoli Enterprise Console 800
- event synchronization
 - installation planning 126
 - removing 940
 - Situation Update Forwarder process 759, 821
 - uninstalling programmatically 940
- event synchronization component
 - changes made to event server 804
 - changing the configuration 822
 - changing the TCP/IP timeout setting 825
 - creating a new rule base 817
 - defining additional monitoring servers 822
 - importing agent class files and rule set 819
 - importing class files and rule set 816
 - installing from the command-line 809
 - Netcool/OMNIBus ObjectServer 743
 - installing with the wizard 805

- event synchronization component (*continued*)
 - Netcool/OMNIBus ObjectServer 740
 - methods of installing 804
 - omegamon.rls file 826
 - overview 13
 - sitconfig.sh command 822
 - sitconfuser.sh command 822
 - uninstalling manually 942
 - upgrading to version 2.3.0.0
 - from the command-line 827
 - using silent installation 829
 - upgrading to version 3.0.0.0
 - from a wizard 826
 - Tivoli Enterprise Console 826
 - using silent installation 813
 - Netcool/OMNIBus ObjectServer 746
- event synchronization configuration fields
 - Netcool/OMNIBus 740
- event synchronization event class
 - attributes 818
- event view 13
- events, forwarding to Netcool/OMNIBus 764
- events, forwarding to Tivoli Enterprise Console 820
- eWAS server 86
- example firewall gateway configuration scenario 873
- explicit rebind 494
- external Web server
 - portal client connection 456

F

- file descriptors limit 137
- file system caching 491
- file, Sentry.baroc 818
- files
 - atr 292
 - cat 292
 - catalog and attribute 275
 - seedkpp.log file 293
 - SQL 275
- Firefox browser
 - tips for using with Tivoli Monitoring 343
- firewall gateway 33
 - activation 869
 - configuration 34, 869
 - configuration for Warehouse Proxy 872
 - determine if you require 22
 - example configuration scenario 873
 - KDE_GATEWAY configuration variable 869
 - uses 868
 - XML configuration document 870
- Firewall gateway 33
- firewall gateway configuration document 870
- firewall support 867
 - editing the partition file
 - UNIX and Linux 867
 - Windows 866

- firewall support (*continued*)
 - sample partition file 867
- firewalls
 - achieving interoperability across 861
 - and Network Address Translation 862
 - factors in implementation 861
 - flow of connection establishment 861
 - implementing with ephemeral pipe 863
 - implementing with firewall gateway 868
 - implementing with partition files 864
 - multiple network interface cards 876
 - NAT 876
 - NIC 876
 - number of internet zones 862
 - options for 861
 - permissions at 861
 - portal server interface, defining 876, 877
 - scenarios 878
 - server address continuity across 862
- firewalls, achieving interoperability across 861
- fix packs 370
- fixes, obtaining 958
- FLUSH PACKAGE CACHE
 - statement 494
- forwarding events to Tivoli Enterprise Console 820
- forwarding situation events
 - to Tivoli Enterprise Console 800
- fragmentation 493
- fsync() 439
- fully qualified path names, using 135

G

- GET SNAPSHOT command 502
- Global Location Broker 657
- Global Zone 135
- globalization, installing support 339
- GROUP BY clause 494
- GSK_V3_CIPHER_SPECS parameter 158

H

- hard disk drive footprint 511
- hardware design 495
- hardware requirements
 - distributed systems 84
 - System z 88
- hash JOINS 497
- health checks 121
 - daily 122
 - monthly 123
 - quarterly 123
 - weekly 123
- heaps 495, 498
- heartbeat 24
 - tracking 24
- heartbeat interval 426
 - CTIRA_HEARTBEAT environment variable 426
 - performance issues 426

- heartbeat interval (*continued*)
 - setting the interval 426
- heartbeat monitoring 426
 - performance issues 426
- high availability and disaster recovery 43
 - agent and remote monitoring server considerations 45
- hub monitoring server
 - considerations 43
 - portal server considerations 44
- Summarization and Pruning Agent
 - considerations 47
- Tivoli Data Warehouse
 - considerations 46
- Tivoli Performance Analyzer
 - considerations 47
- Warehouse Proxy Agent
 - considerations 46
- high speed link 495
- historical data 13
- host variables 494
- HTTP
 - port assignments 423
- HTTP daemon 422
 - KDE_TRANSPORT environment variable for 422
- HTTPS
 - port assignments 423
- HTTPS daemon 422
 - KDE_TRANSPORT environment variable for 422
- hub monitoring server
 - adding application support on Linux 208
 - adding application support on UNIX 208
 - configuring on Linux 202
 - configuring on UNIX 202, 224, 415
 - installing 196
 - installing on Linux 200
 - installing on Windows 196
 - planning worksheet, Linux or UNIX 834
 - planning worksheet, Windows 833

I

- I/O operations, parallel 495
- I/O, minimizing 492, 495
- IBM Infrastructure Management Dashboards for Servers
 - configuring 314
 - installing 306
 - rollback 317
- IBM Java Web Start
 - using to download the desktop client 367
- IBM Redbooks 957
- IBM Support Assistant 957
- IBM Tivoli License Manager 6
- IBM Tivoli Monitoring
 - agent deployment 371
 - component status, monitoring 426
 - configuring components 411
 - heartbeat monitoring 426

- IBM Tivoli Monitoring (*continued*)
 - Java level, required 165
 - Manage Tivoli Enterprise Monitoring Services 411
 - overview 5
 - remote deployment 371
 - removing 929
 - uninstalling 929
 - upgrading to 145
- IBM Tivoli Monitoring 5.x Endpoint 34
- IBM Tivoli Monitoring for Applications: mySAP Agent 156
- IBM Tivoli Monitoring for Databases: DB2 for Linux, UNIX, and Windows Agent 156
- IBM Tivoli Monitoring for Databases: Oracle Agent 156
- IBM Tivoli Monitoring for Databases: Sybase Server Agent 156
- IBM Tivoli Monitoring V5.x Endpoint agent 160
- IBM Tivoli Monitoring V5.x interoperability 160
- IBM Tivoli Monitoring Web Services
 - adding users 464
 - defining hubs 461
 - verifying configuration 467
- IBM Tivoli Monitoring Web Services, configuring 461
- IBM Tivoli Monitoring: UNIX OS Agent 156
- IBM Tivoli Monitoring: Windows OS Agent 156
- iKeyMan 142
- iKeyman utility 454
- implicit rebind 494
- importing event synchronization class files and rule set 816
- In-Place REORG 493
- inbound connection 495
- indexes
 - cardinality 489
 - clustered 489
 - columns 494
 - dropped 494
 - extensive updates 494
 - new 494
 - perfectly ordered 493
 - rebuilt, reorganizing with REORG 493
 - statistics 489, 494
 - synchronized statistics 494
- information to gather before installing 126
- install 97
- installation
 - hardware requirements, distributed systems 84
 - hardware requirements, System z 88
 - information to gather 126
 - Linux considerations 130
 - monitoring server name guidelines 127
 - operating systems, supported 75
 - order of install 128
 - overview of the process 125
 - required information 126

- installation (*continued*)
 - requirements 75
 - software requirements 89
 - UNIX considerations 130
 - Windows considerations 129
- installation steps 193
- installing 101, 193, 196, 211, 226
 - application support 279
 - Automation Server 220
 - desktop client 272
 - globalization support 339
 - language packs 339
 - Linux monitoring server 200, 215
 - monitoring agents 261
 - monitoring agents on Linux 267, 268
 - monitoring agents on UNIX 267, 268
 - monitoring agents on Windows 262
 - planning worksheets 833
 - silent 841
 - Tivoli Enterprise Monitoring
 - Automation Server 219
 - UNIX monitoring server 200, 215
 - Windows monitoring server 196, 211
- installing all components on a single
 - Windows computer 171
- installing application support
 - catalog and attribute files 275
 - seed data 275
 - SQL files 275
- installing into Solaris zones 135
- installing monitoring agent .baroc files on
 - the event server 819
- InstallPresentation.sh 159
- Integrated Service Management
 - Library 955
- integration 13
- integrity, data 492
- interoperability with version 6.3 158
- INTRA_PARALLEL parameter 497
- IP PIPE 21
- IP SPIPE 21
- IPv4 address data 870
- IPv6 communications protocol 127
- ITMAgentsN script 130
- itmcmd agent command 271, 420
- itmcmd config command 851
- itmcmd manage command 412
- itmcmd resp command 852
- itmcmd server command 209, 420
- itmcmd startAgent command 420
- itmcmd stopAgent command 420
- itmcmd support command 202, 209

J

- Java Runtime Environment,
 - required 165
- Java Web Start 342
 - using to download the desktop
 - client 367
- JRE
 - enabling tracing for 367

K

- KBB_RAS1 trace setting 657
- kcirunas.cfg file 130
- KDC_FAMILIES configuration
 - variable 872
- KDC_FAMILIES environment
 - variable 423, 425
- KDC_FAMILIES variable 863
- KDC_PARTITION environment
 - variable 867
- KDE_GATEWAY configuration
 - variable 869
- KDE_TRANSPORT environment
 - variable 423, 425
- KDE_TRANSPORT variable 863
- KDY0012E message 390
- keyword
 - ephemeral 863
- KGLCB_FSYNC_ENABLED
 - parameter 439
- KHD_CNX_POOL_SIZE environment
 - variable 495
- KHDENV Warehouse Proxy Agent
 - configuration file 517, 519
- KinConfig.exe utility 152
- KNTENV file 869
- Kochi fonts
 - installing on SuSE Linux 339
- KPX_WAREHOUSE_LOCATION 657
- KPX_WAREHOUSE_LOCATION
 - variable 863

L

- language packs, installing 339
- large environments 35
- Launchpad 301, 303
- License Manager, IBM Tivoli 6
- link, high speed 495
- Linux 200, 202, 215, 216
 - /etc/hosts file 134
 - adding application support 208
 - backing up a current installation 154
 - configuring 851
 - configuring desktop client 275
 - configuring monitoring agents 269
 - configuring portal desktop client 275
 - configuring the portal server 243, 250
 - EIB tables 923
 - file descriptors limit 137
 - file permissions for monitoring
 - agents 270
 - fully qualified path names 135
 - host names, setting 134
 - installation considerations 130
 - installation user account 134
 - installing as root 134
 - installing monitoring agents 268
 - installing the desktop client from
 - installation media 273
 - installing the portal server 242
 - maxfiles parameter 137
 - monitoring agents 267
 - network interface cards 135
 - NFS environment 135

Linux (*continued*)

- planning 130
- portal server installation 241
- response file configuration 851
- response file installation 849
- silent configuration 851
- silent installation 848, 849
- starting monitoring agents 271
- starting portal server 260
- TCP/IP network services 134
- Tivoli Enterprise Monitoring
 - Automation Server 224, 415
- Linux monitoring server
 - firewall support 867
 - KDC_PARTITION 867
 - NAT 867
 - network address translation
 - (NAT) 867
- LOBs, DB2 491
- Local Location Broker 657
- localhost requirements
 - Linux 440
- location broker 862
- Location Broker
 - Global 657
 - Local 657
- lock list 498
- LOCKLIST parameter 498
- Log Buffer parameter 492
- log retain logging, definition and
 - purpose 492
- LOGBUFSZ parameter 492, 498
- LOGFILSIZ parameter 492, 498
- logging
 - circular 492
 - definition and purpose 492
 - log retain 492
 - mirroring 492
 - performance 492
 - records 492
- LOGPRIMARY parameter 498
- logs
 - for desktop client 367
- lz.config file 869

M

- mainframe 62
 - deployment considerations 62
- maintenance 119, 121
 - installing 370
 - post-upgrade 120
 - pre-upgrade 119
- maintenance utilities for the
 - database 493
- Manage Tivoli Enterprise Monitoring
 - Services 411
 - adding application support 292, 294
 - agents, configuring 417
 - application support, adding 292, 294
 - defining SOAP hubs 461
 - editing the partition file
 - UNIX and Linux 867
 - Windows 866
 - itmcmd manage command 412
 - monitoring agents, configuring 417
 - monitoring server, configuring 412

- Manage Tivoli Enterprise Monitoring Services *(continued)*
 - starting components 420
 - starting on Linux 412
 - starting on UNIX 412
 - starting on Windows 411
 - stopping components 420
- MAX_QUERYDEGREE parameter 497
- maxfiles parameter 137
- MAXLOCKS parameter 498
- medium environments 35
- memory
 - accessing 492
 - allocation 495
 - buffer pools 492
 - DB2 model 495
 - organization 495
 - parameters 495
 - tuning 495
- memory requirements 85
- Microsoft SQL data warehouse
 - ODBC connection 602
- Microsoft Windows Terminal Services 130
- migrate-export process 453
- migrate-import process 453
- migrated information 152
- MINCOMMIT parameter 492
- mirroring log files 492
- monitoring agent
 - planning worksheet, Linux or UNIX 838
 - planning worksheet, Windows 838
 - removing through the portal 937
 - uninstalling on Linux 935
 - uninstalling on UNIX 935
 - uninstalling on Windows 934
- monitoring agent .baroc files, installing on the event server 819
- monitoring agents
 - configuring 417
 - configuring on Linux 269
 - configuring on UNIX 269
 - definition 11
 - deploying 371
 - deploying OS agents 375
 - file permissions, changing 270
 - installing 261
 - installing on Linux 267, 268
 - installing on UNIX 268
 - installing on Windows 262
 - monitoring agents on Linux 267
 - product codes 883
 - starting 420
 - stopping 420
- monitoring agents, base 5
- Monitoring components 23, 129
 - Agent Builder 33
 - Firewall gateway 33
 - IBM Tivoli Monitoring 5.x
 - Endpoint 34
 - Netcool/OMNIBus integration 31
 - Summarization and Pruning
 - Agent 30
 - Tivoli Data Warehouse 30
 - Tivoli Enterprise Console
 - integration 31

- Monitoring components *(continued)*
 - Tivoli Enterprise Monitoring
 - Agent 28
 - Tivoli Enterprise Monitoring Server 24
 - Tivoli Enterprise Portal client 27
 - Tivoli Enterprise Portal Server 25
 - Warehouse Proxy Agent 29
- monitoring environment
 - Tivoli Enterprise Console 803
- monitoring server
 - adding application support on
 - Linux 161, 208, 284
 - adding application support on
 - UNIX 161, 208, 284
 - adding application support on
 - Windows 161, 281
 - agent depot, location 374
 - agent depot, populating 371
 - communications protocol planning
 - worksheet 839
 - configuration file 197, 211
 - configuring 412, 419
 - configuring on Linux 202, 216
 - configuring on UNIX 202, 216
 - DEPOTHOME environment
 - variable 374
 - environment file 197, 211
 - file descriptor limit on UNIX 137
 - forwarding events to Tivoli Enterprise Console 820
 - heartbeat interval, configuring 426
 - installing on Linux 200, 215
 - installing on UNIX 200, 215
 - installing on Windows 196, 211
 - KBBENV configuration file 197, 211
 - naming 127
 - pure events 419
 - starting 209, 420
 - stopping 209, 420
 - Tivoli Event Integration Facility 764, 820
- monthly health checks 123
- motherboard 495
- MPP (Massively Parallel Processors) 495
- multiple network interface cards 135, 876
 - portal server interface, defining 876, 877
- multiple nodes 495
- multiprocessor systems 84

N

- naming conventions 99
- naming the monitoring server 127
- NAT 867, 876
- Netcool SSM agents
 - deployment of 386
- Netcool/OMNIBus 697
 - verifying installation of event
 - synchronization component 765
- Netcool/OMNIBus event synchronization
 - configuration fields 740
- Netcool/OMNIBus integration 31

- Netcool/OMNIBus ObjectServer
 - configuring for program execution
 - from scripts 763
 - configuring the EIF probe 760
 - configuring the error event flow 773
 - updating the database schema 753
- network address translation (NAT) 867, 876
 - portal server interface, defining 876, 877
- network interface cards
 - specifying 135
- network, monitoring bandwidth 495
- new in this release 1
- NFS environment, installing into 135
 - required permissions 135
- NIC 876
 - portal server interface, defining 876, 877
- NOCACHE option 491
- node ID 296
- nodes, multiple 495
- non-NIS Solaris monitoring serve,
 - configuring permissions for 439
- non-OS agents
 - deploying 377
- non-sequential physical data pages 493
- notices 963
- NUM_IOCLEANERS parameter 498
- NUM_IOSERVERS parameter 498

O

- ODBC 29
- ODBC data source
 - uninstalling 938
- OLTP ratio 495
- om_tec.config event-forwarding control file 764, 821
- OMEGAMON
 - uninstalling agents 935
- omegamon.baroc class file 804
- omegamon.rls rule set file 804
- ONLY ON KEY COLUMNS clause 494
- Open Database Connectivity 29
- operating systems
 - no longer supported 145
- operating systems, supported 75
- optimizer
 - access plan creator 489
 - intra-partition parallelism
 - determination 498
 - purpose 494
 - statement binding with dynamic SQL 494
 - statistics influence 489
- Optional Primary Network Name
 - option 135
- options
 - Security: Validate User 139
- Oracle data warehouse
 - ODBC connection 627
- OS Cluster 44
- overview 5

P

- packages
 - cache for dynamic and static SQL 498
 - invalid 494
 - static SQL 494
- page cleaners, asynchronous 498
- paging and swapping 495
- parallel I/O operations 495
- partition file
 - editing
 - UNIX and Linux 867
 - Windows 866
 - sample 867
- partition files
 - implementing 864
 - sample scenarios 865
- password authentication, configuring 139
- password encryption 137
- PCKCACHESZ parameter 498
- Performance Analyzer
 - Configuring historical data collection 432
 - Silent configuration 431
- Performance tuning 469
- performance tuning for DB2 databases 489
- physical data pages, non-sequential 493
- placing components 23
- planning 19
 - hardware requirements, distributed systems 84
 - hardware requirements, System z 88
 - installation 125
 - Linux considerations 130
 - operating systems, supported 75
 - requirements 75
 - software requirements 89
 - Tivoli Data Warehouse 510
 - Tivoli Enterprise Console 803
 - UNIX considerations 130
 - Windows considerations 129
 - worksheets 833
- planning, upgrade 145
- pooling, connection 495
- populate_agents.sql script 271
- port number assignments
 - configuring client connections for the portal server 421
 - configuring for monitoring agents 423
 - configuring for the monitoring server 421
 - controlling 421
- portal browser client
 - starting 366, 367
- portal client
 - connecting to an external Web server 456
 - external Web server, connecting to 456
- Portal client
 - browser client 11
 - defined 11
 - client 11
- Portal client (*continued*)
 - components
 - browser client 11
 - desktop client 11
 - desktop client 11
 - defined 11
- portal desktop client
 - adding application support on Linux 291
 - adding application support on Windows 290
 - configuring on Linux 275
 - downloading with IBM Java Web Start 367
 - installing on Linux from installation media 273
 - installing on Windows from installation media 272
 - planning worksheet, Linux 839
 - planning worksheet, Windows 838
 - starting 366
- portal server
 - adding application support on Linux 288
 - adding application support on Windows 287
 - backing up 240, 261
 - configuring on Linux or AIX 243, 250
 - installing 226
 - installing on Linux or AIX 241, 242
 - installing on Windows 227
 - restoring 240, 261
 - starting on Linux 260
- portal server database
 - backing up 155
- portal server interface
 - defining on Linux or UNIX 877
 - defining on Windows 876
- post-deployment 119
 - applying maintenance 119, 370
 - maintenance 121
- post-installation 101
- post-upgrade 120
- pre-deployment 19
 - accelerating 65
 - additional ports 34
 - agent deployments 47
 - firewall gateway determination 22
 - high availability and disaster recovery 43
 - Monitoring component placement 23
 - multi-hub environments 63
 - planning checklist 19
 - project management and planning 66
 - remote deployment 36
 - sizing your Tivoli Monitoring hardware 34
 - Software Product Compatibility Reports 43
 - staffing 72
 - task estimates 66
 - understanding your network 21
- pre-installation 97
- prefetch
 - definition 489
 - list 489
- prefetch (*continued*)
 - list sequential 489
 - quantity 494
 - sequential 489
- prerequisite scanning
 - location 49
 - monitoring agents 48
 - remote 51
 - running 49
 - sample report 51
 - stand-alone 49
- prerequisites
 - for upgrading 146
 - hardware 75
 - software 75
- privacy policy 965
- problem determination and resolution 960
- problem resolution 957
- Problems
 - submitting 960
- processor configurations 495
- processor requirements 84
- proddsc.tbl 137, 850
- product codes 883
 - Linux 137
 - UNIX 137
- product documentation 950
- product overview 5
- project management and planning 66
- public/private key pair 454
- publication tool, schema 525

Q

- quarterly health checks 123
- queries 494
- queue delays, causes 495

R

- RAID device 491
- ratios, scaling up 495
- raw data size, calculating 495
- rc.itm script 130
- read operations, excessive 493
- REBIND utility
 - place in order of utilities 494
 - purpose 494
 - SQL types supported 494
- rebinding, explicit and implicit 494
- recovery, roll-forward 492
- Redbooks 489, 950, 955, 957
- REDISTRIBUTE DATABASE PARTITION GROUP utility 494
- registry variables 491, 501
- release information 1
- remote deployment 36, 371
 - deploying OS agents 375
 - managing agent depot 374
 - OS agents, deploying 375
 - populating agent depot 371
 - sharing an agent depot 375
 - tacmd createNode command 375
- Remote Execution and Access 115

- remote monitoring server
 - configuring on Linux 216
 - configuring on UNIX 216
 - installing 211
 - installing on Linux 215
 - installing on UNIX 215
 - installing on Windows 211
 - planning worksheet, Linux or UNIX 836
 - planning worksheet, Windows 835
- removing 929
 - agents 937
 - component 934
 - Linux 935
 - UNIX 935
 - Warehouse Proxy 937
 - Windows 934
 - event synchronization 940
 - from the portal 937
 - monitoring component 934
 - Linux 935
 - UNIX 935
 - Warehouse Proxy 937
 - Windows 934
 - monitoring environment 929
 - Linux 933
 - UNIX 933
 - Windows 929
 - OMEGAMON Monitoring Agent 935
- REORG utility
 - classic 493
 - In-Place 493
 - types 493
- reporting software 13
- required hardware
 - distributed systems 84
 - System z 88
- required information for installation 126
- required software 89
- requirements
 - disk storage 85
 - hardware 75
 - memory 85
 - processor 84
 - software 75
- resources 949
 - education 949
 - education offerings 950
 - IBM Tivoli Monitoring 6 Welcome Kit 949
 - other 951
 - product documentation 950
 - Redbooks 950
 - service offerings 951
 - support Web sites 949
- response file
 - use with schema publication tool 525
- response time, delayed 495
- roll-forward recovery 492
- rule base
 - creating new 817
 - importing existing rule base into 817
 - modifying 818
- RUNSTATS utility
 - formats 494
 - list of purposes 494
 - options 494

- RUNSTATS utility (*continued*)
 - precaution 494
- RXA 115
 - Disabling User Account Control 116
 - Windows 115

S

- SAMPLE DETAILED clause 494
- scalability 495
- scaling up ratios 495
- scenarios
 - partition files 865
 - secureMain 926
- schema publication tool 525
 - updated mode 527
- Schema Publication Tool
 - database compression 529
- scripts
 - autostart 130
- secure protocols
 - Internet Inter-ORB Protocol (IIOP) 138
 - Interoperable Object Reference (IOR) protocol 138
 - Secure Hypertext Transport Protocol (HTTPS) 138
- secureMain utility 925
 - examples 926
 - scenarios 926
 - usage 925
- securing a Linux or UNIX installation 925
- securing communication between components 138
- securing the SOAP server 142
- security
 - firewall scenarios 878
 - HTTPS 454
 - Secure Hypertext Transport Protocol (HTTPS) 454
 - SSL 454
- security options
 - for communications 137, 138
 - for SOAP Server 137
 - for user authentication 137
 - single sign-on 141
- Security: Validate User option 139
- seed data 275
- seeding 275
- seedkpp.log file 293
- seek time 495
- self-describing agents 48
 - configuring seeding 206
 - dynamically controlling 205
 - installation 394
 - manage installations 394
- self-describing monitoring agents 69
- Sentry.baroc class file 804
- Sentry.baroc file 818
- Sentry2_0_Base class 126, 804
- sequential prefetch 489
- Service Console, CT 422
- Service Management Connect 955
- SET RUNTIME DEGREE command 498
- setting, KBB_RAS1 trace 657
- sharing 63

- sharing (*continued*)
 - customization 63
- SHEAPTHRES parameter 497
- shortcut for launching desktop client 369
- silent installation 841
 - Linux 848, 849
 - UNIX 848, 849
 - Windows 841
 - from command-line 844
 - using SMS 845
- silent response file 405
- silent-installation files 843, 852
- single sign-on capability 141
- single-computer installation
 - order 128
- sitconfig.sh command 822
- sitconfuser.sh command 822
- sitpad tool 167
- Situation Event Console 13
- Situation Update Forwarder
 - process 759, 821
- size of raw data, calculating 495
- small environments 35
- Small Local Zone 135
- Small Local Zone with /opt directory share 135
- SMC 955
- SMP (Symmetric Multiprocessor) 495
- SOAP server
 - adding users 464
 - defining hubs 461
 - verifying configuration 467
- SOAP Server
 - security options for 137
- SOAP server, configuring 461
- SOAP server, securing 142
- sockets 495
- software requirements 89
- Software Support
 - contacting 959
 - overview 957
 - receiving weekly updates 958
- Solaris zones
 - installing into 135
 - types of 135
- Solution Installer 843
- SORTHEAP parameter 491, 497, 498
- SQL
 - dynamic 494
 - optimizer 494
 - static 494
 - types 494
- SQL files 275
- SQL query 663
 - assigning to a workspace 665
 - creating 663
 - WAREHOUSELOG 665
- SQL Server, Microsoft 82
- SQL statements
 - schema publication tool for generating 525
- SSO 141
- staffing 72
- Star JOINs 497

- starting Manage Tivoli Enterprise Monitoring Services
 - itmcmd manage command 412
 - on Linux 412
 - on UNIX 412
 - on Windows 411
- starting monitoring agents
 - itmcmd agent command 271
- starting the monitoring server
 - itmcmd server command 209
- starting the portal server 260
- Startup Center 103
- stash file 142
- statistics
 - collected for RUNSTATS 494
 - comparing current and previous 494
 - distribution 494
 - index 494
 - queries 494
 - types of collections 494
- stopping the monitoring server
 - itmcmd server command 209
- Summarization and Pruning Agent 30
 - configuration for autonomous operation 658
 - configuring communications for DB2
 - for Linux, UNIX and Windows 562
 - configuring JDBC connection 643
 - considerations 39
 - deployment task estimates 68
 - description of 519
 - high availability and disaster recovery
 - considerations 47
 - installing on DB2 for Linux, UNIX and Windows 562
 - product code 883
 - range partitioning 533
 - starting 656
- support assistant 957
- support Web sites 949
- supported databases
 - Tivoli Data Warehouse 82
- supported operating systems 75
- swapping 495
- switches 501
- sysadmin user account 139
- system caching, files 491
- System Managed Space (SMS) table space 490
- System Monitor Agents 399
 - configuring
 - on Linux or UNIX 406
 - on Windows 402
 - installing 404
 - on Linux or UNIX 404
 - on Windows 400
 - on Windows 64-bit processors 401
 - uninstalling
 - on Linux or UNIX 408
 - on Windows 403
- System z
 - required hardware 88
- systems, multiprocessor 84

T

- table spaces
 - containers 490
 - definition and purpose 490
 - multiple page sizes 492
 - types 490
- tacmd addBundles command 265, 373, 379, 380, 386
- tacmd addSystem command 386, 390
- tacmd AddSystem command 265
- tacmd clearDeployStatus command 390
- tacmd commands
 - remote deployment 390
- tacmd configureSystem command 381, 386, 390
- tacmd createNode command 375, 379, 386, 390
 - requirements 376
 - using 377
- tacmd executecommand command
 - installing application support for an IBM Tivoli Composite Application Manager agent 275
- tacmd getDeployStatus command 384, 390
- tacmd getFile command
 - installing application support for an IBM Tivoli Composite Application Manager agent 275
- tacmd putfile command
 - installing application support for an IBM Tivoli Composite Application Manager agent 275
- tacmd removeSystem command 380, 390
- tacmd restartAgent command 380
- tacmd startAgent command 380
- tacmd stopAgent command 380
- tacmd updateAgent command 380, 386, 390
- tacmd viewDepot command 379
- taudit tool 167
- TCP 21
- TCP/IP timeout setting, changing on event server 825
- TDS 137
- TEC_Generic class 818
- tec.baroc file 818
- Technotes 955
- TEMS name 296
- TEPS/e 14
- tepsBackup command 240, 261
- tepsRestore command 240, 261
- Terminal Services 130
- throughput 489, 495
- Tivoli Common Reporting 693
- Tivoli Data Warehouse 30
 - backing up the database 156
 - creating a Windows user 544
 - creating database objects for 525
 - DB2 for Linux, UNIX, and Windows
 - database, increasing the size of 514
 - disk requirements, understanding 513
 - estimating database size 510
 - high availability and disaster recovery
 - considerations 46
 - planning 510

- Tivoli Data Warehouse (*continued*)
 - range partitioning 531
 - sharing 63
 - supported databases 82
 - tables and indices 537
 - uninstalling the Warehouse
 - Proxy 937
 - upgrading 156
 - Windows user, creating 544
 - worksheet 512

- Tivoli Data Warehouse database
 - backing up 156

- Tivoli Directory Server 137

- Tivoli Enterprise Console
 - changing the TCP/IP timeout
 - setting 825
 - determining when to use 803
 - event forwarding 799
 - event integration 799
 - forwarding events to 820
 - information required for event forwarding 126
 - integration 800
 - planning 800

- Tivoli Event Integration Facility 820

- Tivoli Enterprise Console integration 31
 - deployment task estimates 68

- Tivoli Enterprise Monitoring Agent 11, 28
 - deployment task estimates 69

- Tivoli Enterprise Monitoring Agents
 - configuring port number
 - assignments 423

- Tivoli Enterprise Monitoring Automation Server
 - configuring on Linux 224, 415
 - configuring on UNIX 224, 415
 - installing 219
 - planning worksheet, Windows 834, 835

- Tivoli Enterprise Monitoring Server 10, 24
 - configuring port number 421

- high availability and disaster recovery
 - considerations 43
- hub considerations 35
- multiple hub considerations 35
- remote considerations 35
- remote high availability and disaster recovery
 - considerations 45

- Tivoli Enterprise Portal 11
 - authorizing users 139
 - new in this release 1

- Tivoli Enterprise Portal client 27
 - considerations 40

- Tivoli Enterprise Portal Server 25
 - configuring port number assignments
 - for client connections 421
 - considerations 37
 - high availability and disaster recovery
 - considerations 44

- Tivoli Enterprise Portal Server extended services (TEPS/e)
 - definition of 14

- Tivoli Event Integration Facility 764, 820
- Tivoli Management Services 5

- Tivoli Monitoring
 - components 6
- Tivoli Monitoring environment
 - additional ports used 34
 - components 23
 - maintenance 370
 - sizing your hardware 34
- Tivoli Performance Analyze 40
- Tivoli Performance Analyzer
 - high availability and disaster recovery considerations 47
- tools
 - benchmarking 501
 - db2batch 501
 - Enable Multipage File Allocation 491
 - Event Monitor 501
 - Explain Facility 501
 - monitoring 501
 - Snapshot Monitor 501
- topology view 149
- trace setting, KBB_RAS1 657
- trademarks 965
- transactions 492
- troubleshooting 957
- tuning
 - DB2 performance on Windows, UNIX, and Linux 489
 - guidelines 497
- tuning parameters
 - AGENTPRI 497
 - ASLHEAPSZ 497
 - CHNGPGS_THRESH 498
 - DBHEAP 498
 - DFT_DEGREE 498
 - INTRA_PARALLEL 497
 - LOCKLIST 498
 - LOGBUFSZ 498
 - LOGFILSIZ 498
 - LOGPRIMARY 498
 - MAX_QUERYDEGREE 497
 - MAXLOCKS 498
 - NUM_IOCLEANERS 498
 - NUM_IOSERVERS 498
 - PCKCACHESZ 498
 - SHEAPTHRES 497
 - SORTHEAP 498
- type 883

U

- ulimit values 440
- Uni-Processor 495
- uninstallation 929
 - agents 937
 - component 934
 - Linux 935
 - UNIX 935
 - Warehouse Proxy 937
 - Windows 934
 - from the portal 937
 - monitoring component 934
 - Linux 935
 - UNIX 935
 - Warehouse Proxy 937
 - Windows 934
 - monitoring environment 929
 - Linux 933

- uninstallation (*continued*)
 - monitoring environment (*continued*)
 - UNIX 933
 - Windows 929
 - Warehouse Proxy 937
- uninstalling
 - event synchronization 940
 - event synchronization component manually 942
 - IBM Tivoli Monitoring
 - environment 929
 - ODBC data source 938
 - OMEGAMON Monitoring Agents 935
 - Warehouse Proxy Agent 937
 - UNIX 200, 202, 215, 216
 - /etc/hosts file 134
 - adding application support 208
 - backing up a current installation 154
 - configuring 851
 - configuring monitoring agents 269
 - EIB tables 923
 - existing installation
 - installing into 137
 - file descriptors limit 137
 - file permissions for configuring monitoring agents 270
 - fully qualified path names 135
 - host names, setting 134
 - installation considerations 130
 - installation user account 134
 - installing as root 134
 - installing monitoring agents 268
 - maxfiles parameter 137
 - monitoring agents 267
 - monitoring server parameter 439
 - network interface cards 135
 - NFS environment 135
 - planning 130
 - response file configuration 851
 - response file installation 849
 - silent configuration 851
 - silent installation 848, 849
 - Solaris zones
 - installing into 135
 - starting monitoring agents 271
 - TCP/IP network services 134
 - Tivoli Enterprise Monitoring Automation Server 224, 415
 - UNIX environment 67
 - UNIX monitoring serve
 - configuring permissions for non-NIS Solaris 439
 - non-NIS Solaris monitoring serve 439
 - UNIX monitoring server
 - firewall support 867
 - KDC_PARTITION 867
 - NAT 867
 - network address translation (NAT) 867
 - unsupported operating systems 145
 - updated mode
 - schema publication tool 527
 - updating the Netcool/OMNIbus database schema 753
 - upgrade 119

- upgrade (*continued*)
 - order 128
 - steps 119
- upgrade planning 145
- upgrading
 - deciding which components 149
 - from IBM Tivoli Monitoring V6.1 161
 - overview 161
 - from IBM Tivoli Monitoring V6.2 161
 - overview 161
 - Java level, required 165
 - migrated information 152
 - order 128
 - required order 150
 - Tivoli Data Warehouse 156
- upgrading the Tivoli Data Warehouse 156
- upgrading to IBM Tivoli Monitoring 145
- upgrading to V6.2
 - required order 150
- upgrading to version 6.2/6.3
 - planning 145
- user accounts 139
- user authentication
 - options for 137
- user authority when installing on Linux or UNIX 134
- using Java Web Start to download the desktop client 368
- utilities
 - KinConfig.exe 152
- utilization, CPU 495
- ux.config file 869
- ux.ini file 869

V

- variables
 - KDC_FAMILIES 863
 - KDE_TRANSPORT 863
 - KPX_WAREHOUSE_LOCATION 863
- variables, host 494
- view 13
 - topology 149
- VMWare ESX Server 43

W

- Warehouse load projection
 - spreadsheet 37
- Warehouse Proxy
 - removing 937
 - tuning performance 666
 - uninstalling 937
- Warehouse Proxy Agent 13, 29
 - 32-bit 519, 521
 - 64-bit 519, 521, 522
 - installation process 519
 - configuration 656
 - configuration for autonomous operation 658
 - port assignment for 658
 - considerations 37
 - deployment task estimates 68
 - file descriptor limit on UNIX 137

- Warehouse Proxy Agent *(continued)*
 - high availability and disaster recovery considerations 46
 - installing 656
 - product code 883
 - range partitioning 532
 - support for multiple 656
 - Warehouse load projection spreadsheet 37
- WAREHOUSEAGGREGLOG table 689
 - ENDTIME column 689
 - LOGTMZDIFF column 689
 - MAXWRITETIME column 689
 - MINWRITETIME column 689
 - OBJECT column 689
 - ORIGINNODE column 689
 - ROWSREAD column 689
 - STARTTIME column 689
- WAREHOUSELOG query 665
- WAREHOUSELOG table 689
 - ENDQUEUE column 689
 - ENDTIME column 689
 - ERRORMSG column 689
 - EXPORTTIME column 689
 - OBJECT column 689
 - ORIGINNODE column 689
 - ROWSINSERTED column 689
 - ROWSRECEIVED column 689
 - ROWSSKIPPED column 689
 - STARTEXPART column 689
 - STARTQUEUE column 689
 - STARTTIME column 689
 - WPSYSNAME column 689
- WAREHOUSELOG workspace 665
- Web browser
 - changing locations on UNIX or Linux 364
- Web Health Console, IBM Tivoli Monitoring V5.1
 - Tivoli Enterprise Portal visualization for 160
- Web Start
 - launching from IBM Java Control Panel 368
 - launching the desktop client from a browser 368
 - launching the desktop client from the command-line 368
- weekly health checks 123
- what is new 1
- WHERE clause 494
- Windows 196, 211
 - backing up a current installation 152
 - installing the desktop client 272
 - monitoring agents 262
 - monitoring server
 - configuration file 227, 272
 - environment file 227, 272
 - KBBENV configuration file 227, 272
 - response file 841
 - silent installation 841
- Windows environment 67
- Windows Terminal Services 130
- work
 - roll-forward recovery 492
 - rollbacks 492
- worksheets 833
 - communications protocol 839
 - desktop client on Linux 839
 - desktop client on Windows 838
 - Linux or UNIX hub monitoring server 834
 - Linux or UNIX remote monitoring server 836
 - Linux or UNIX Tivoli Enterprise Monitoring Automation Server 835
 - Linux portal server 837
 - monitoring agent on Linux or UNIX 838
 - monitoring agent on Windows 838
 - monitoring server communications protocol 839
 - portal desktop client on Linux 839
 - portal desktop client on
 - Windows 838
 - Windows hub monitoring server 833
 - Windows portal server 836
 - Windows remote monitoring server 835
 - Windows Tivoli Enterprise Monitoring Automation Server 834
- workspace, WAREHOUSELOG 665
- workspaces, creating 665
- world permission
 - required for NFS access 135

Z

- z/OS server 67



Printed in USA

SC22-5445-01

