

Tivoli® OMEGAMON XE on z/VM and Linux
Version 4.2.0

Planning and Configuration Guide



Tivoli® OMEGAMON XE on z/VM and Linux
Version 4.2.0

Planning and Configuration Guide



Note

Before using this information and the product it supports, read the information in “Notices” on page 111.

This edition applies to Version 4, Release 2, Modification 0 of OMEGAMON XE on z/VM and Linux (product number 5698-A36) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 2006, 2010.

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

Contents

Figures	vii
Tables.	ix
Preface	xi
Intended audience	xi
What this guide contains	xi
<hr/>	
Part 1. Planning your configuration	1
Chapter 1. Overview of the OMEGAMON XE on z/VM and Linux monitoring agent	3
Tivoli Enterprise Portal	4
What's new in Version 4.2.0	5
Replacement of the z/VM Systems: default workspace with the z/VM Systems: System Health workspace.	5
Overall system view	5
New product-provided situations	6
Tivoli Management Services components	6
IBM Tivoli Monitoring features	7
Dynamic linking to cross-product workspaces	9
OMEGAMON XE on z/VM and Linux monitoring agent architecture overview	9
IBM Tivoli OMEGAMON XE zSeries products	10
Serviceability	12
Standards supported	12
Interoperability with other products	13
Chapter 2. Planning your OMEGAMON XE on z/VM and Linux configuration	15
Understanding and designing your configuration	16
Tivoli Enterprise Portal and Tivoli Enterprise Portal Server	16
Tivoli Enterprise Monitoring Servers - hub and remote	17
Tivoli Enterprise Console.	19
Tivoli Data Warehouse	19
Warehouse Proxy planning	19
Chapter 3. Installation preparation	21
Software and hardware prerequisites	21
Required software	21
Supported hardware	25
Product packaging	25
The IBM Tivoli OMEGAMON XE on z/VM and Linux product package	26
Installation Flow	29
About VMSES/E	29
Installation flow	29
Chapter 4. Upgrading IBM Tivoli Monitoring.	33
Upgrading an existing IBM Tivoli Monitoring environment	33
Installation in a new environment.	33
<hr/>	
Part 2. Configuration required for this monitoring agent	35
Chapter 5. Configuration required for the OMEGAMON XE on z/VM and Linux monitoring agent	37
Required order of tasks for viewing data at the monitoring agent	38

Configuration steps	40
Step 1. Enabling the collection of data	40
Configuring TCP/IP on z/VM	40
Enabling the CP Monitor domains	40
Step 2. Estimating the size of the PERFOUT DCSS	42
Step 3. Defining a DCSS on z/VM	45
Helpful tips when defining your DCSS on z/VM	45
Running out of memory in the DCSS	46
Step 4. Issue the FC MONCOLL SEGOUT ON PERFOUT command	47
Step 5. Configuring the DCSS device driver on the Linux guest	47
Determining the start and end addresses of the PERFOUT DCSS	47
Step 6. Loading the DCSS device driver	48
Step 7. Adding the PERFOUT DCSS to your Linux guest	49
Step 8. Loading the PERFOUT DCSS at startup time	49
DCSS naming scheme	49
Step 9. Enabling the collection of Linux data	49
Starting data collection for User ApplData at the Linux guest	50
Enable the collection of Linux data automatically at startup time	51
Step 10. Enabling dynamic workspace linking	52
Step 11. Enabling Take Action commands (optional)	53
z/VM requirements for Take Action commands	53
Setting the environment variables during monitoring agent installation	54
Linux guest requirements for Take Action commands	54
Guidelines for issuing Take Action commands	56
Step 12. Installing the Language Packs (optional)	57
Starting the monitoring agent	58
Stopping the monitoring agent	59
Chapter 6. Defining user IDs and security	61
Defining the list of user IDs for the Command Processor	61
Defining user IDs in IBM Tivoli Monitoring	61
Part 3. Completing your configuration	63
Chapter 7. Performance and storage considerations	65
Understanding how data are collected	65
Determining which systems to monitor	66
Understanding historical data	66
Determining which types of historical data to collect	67
Disk capacity planning for historical data	68
Defining and running situations	69
Designing workspaces	70
Reducing the number of rows retrieved	70
Reducing the number of attributes retrieved	70
Applying the same query to multiple views in a workspace	71
Adjusting the auto-refresh rate	71
Chapter 8. Serviceability	73
Log files	73
Definitions of variables for RAS1 logs	74
Tivoli Enterprise Monitoring Server on Windows computers or on UNIX computers	75
Tivoli Enterprise Portal	75
Tivoli Enterprise Portal Server	75
Tivoli Data Warehouse and the warehouse proxy	75

Part 4. Appendixes.	77
Planning worksheets	79
Worksheet: Your overall configuration	79
Worksheet: Your monitoring agent configuration	80
Worksheet: Planning communication protocols for the monitoring agent when the monitoring server is on a distributed system	81
Worksheets: Information to gather when configuring your portal server on Windows or Linux	82
Worksheet: Information for configuring your portal server on Windows	82
Worksheet: Information for configuring your portal server on Linux	82
Worksheet: Planning the communication protocols for the portal server.	83
Worksheets: Information to gather when configuring your monitoring server on a distributed system	84
Worksheet: Information for configuring your hub monitoring server on a distributed system	84
Worksheet: Planning communication protocols for the monitoring server on a distributed system	85
Worksheets: Information to gather when putting your hub monitoring server on a z/OS system	86
Worksheet for configuring your hub monitoring server on a z/OS system	86
Worksheet for configuring your communications protocols for a hub monitoring server on a z/OS system	88
Worksheets: Information to gather when configuring your portal desktop client on Windows or on Linux	91
Worksheet: Information for configuring your portal desktop client on Windows	91
Worksheet: Information for configuring your portal desktop client on Linux	91
Worksheet: Planning the communication protocols for the portal desktop client on a Windows system	92
Specifying communication protocols between components	93
Finding the information you need.	95
Planning tasks	95
Understanding temporary defects, limitations, and workarounds	96
Upgrading from an earlier release	96
Installing and uninstalling	97
Configuring.	97
Tuning	98
Administration.	98
Diagnosis	99
Using	99
Support for problem solving	101
Using IBM Support Assistant	101
Obtaining fixes	101
Receiving weekly support updates	102
Contacting IBM Software Support	102
Determining the business impact	103
Describing problems and gathering information	103
Submitting problems	104
Documentation library.	105
OMEGAMON XE on z/VM and Linux library	105
OMEGAMON XE on z/VM and Linux online help	106
IBM Tivoli Monitoring publications	106
Tivoli Enterprise Portal help system	107
Related publications	107
Accessing terminology online.	108
Accessing publications online	108
Ordering publications.	109
Accessibility	109
Notices	111

Trademarks	113
Glossary	115
Index	119

Figures

1.	Tivoli Enterprise Portal sample workspace for OMEGAMON XE on z/VM and Linux	4
2.	OMEGAMON XE on z/VM and Linux architecture	10
3.	IBM Tivoli Monitoring components	16
4.	IBM Tivoli OMEGAMON XE on z/VM and Linux product packaging	27
5.	Installation and configuration overview	30

Tables

1.	New predefined situations	6
2.	IBM Tivoli OMEGAMON XE zSeries products	11
3.	Supported versions of the z/VM operating system	22
4.	Required order of tasks for viewing data at the monitoring agent	38
5.	CP Monitor domains that must be enabled	41
6.	Capacity planning for historical data	68
7.	Locations of various types of logs.	73
8.	Worksheet for designing your overall configuration	79
9.	Worksheet for configuring your monitoring agent	80
10.	Worksheet for specifying communication protocols for the monitoring agent when the monitoring server is on a distributed system	81
11.	Worksheet for configuring your portal server on Windows	82
12.	Worksheet for configuring your portal server on Linux	82
13.	Worksheet for specifying communication protocols for the portal server	83
14.	Worksheet for configuring your hub monitoring server on a distributed system	84
15.	Worksheet for specifying communication protocols for a monitoring server on a distributed system	86
16.	Worksheet for configuring the hub monitoring server on z/OS	86
17.	Worksheet for configuring the communications protocols for a hub monitoring server on z/OS	88
18.	Worksheet for configuring your portal desktop client on Windows	91
19.	Worksheet for configuring your portal desktop client on Linux	91
20.	Worksheet for specifying communication protocols for the portal desktop client	92
21.	Summary: Planning communication protocols for IBM Tivoli Monitoring components	93
22.	Planning tasks for IBM Tivoli Monitoring and monitoring agents	95
23.	Changes to product levels since the installation media was created for the IBM Tivoli Monitoring and monitoring agents	96
24.	Upgrade tasks for IBM Tivoli Monitoring and the monitoring agents	96
25.	Installing tasks for the IBM Tivoli Monitoring and monitoring agents	97
26.	Configuration tasks for the IBM Tivoli Monitoring and monitoring agents	97
27.	Tuning the IBM Tivoli Monitoring and monitoring agents	98
28.	Administrative tasks for the IBM Tivoli Monitoring and monitoring agents	98
29.	Diagnosis tasks for the IBM Tivoli Monitoring and monitoring agents	99
30.	Use information for the IBM Tivoli Monitoring and monitoring agents	99

Preface

IBM® Tivoli® OMEGAMON® XE on z/VM® and Linux® version 4 release 2 (V4.2.0) is composed of the z/VM agent and the Linux OS agent. Both agents are members of the IBM Tivoli Management Services family, a suite of products that monitors and manages systems and network applications on a variety of operating systems and provides workstation-based reports you can use to track trends and troubleshoot system problems.

The OMEGAMON XE on z/VM and Linux product collects system metrics and performance information about the z/VM operating system and Linux on IBM System z® guest systems. The data collected are displayed by the Tivoli Enterprise Portal application. See the Tivoli Enterprise Portal online help and publications for details on this interface.

This publication describes how to plan your deployment of the OMEGAMON XE on z/VM and Linux software, and how to configure the software in your environment.

To install and to set up the software, use the *IBM Tivoli OMEGAMON XE on z/VM and Linux Program Directory* for this monitoring agent, along with the *IBM Tivoli Monitoring Installation and Setup Guide*. You also use this guide for certain configuration steps required by this monitoring agent. Use the information in the *IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide* to monitor z/VM and to monitor Linux on System z.

Intended audience

The primary audience for this publication is the z/VM system programmer or system administrator who is responsible for ensuring the availability of z/VM system resources. This publication is also addressed to the Linux on System z programmer or system administrator who is responsible for ensuring the availability of systems running in Linux guests.

The system programmer or system administrator responsibilities include:

- Planning for and overseeing product installation
- Troubleshooting system and performance problems
- Analyzing performance data for problem determination
- Providing historical performance data for trend analysis

Users of this publication should be familiar with the following topics:

- Performance monitoring concepts
- IBM Tivoli Monitoring and the Tivoli Enterprise Portal interface
- The Linux on System z® operating system and its associated concepts
- The z/VM operating system and its associated concepts
- The Performance Toolkit for VM
- The Microsoft® Windows® operating system

What this guide contains

This guide contains the following sections:

Part 1 - Planning your configuration

Part 1 contains the following chapters:

- Chapter 1, “Overview of the OMEGAMON XE on z/VM and Linux monitoring agent,” on page 3 introduces the OMEGAMON XE on z/VM and Linux monitoring agent. This chapter also introduces the features and components of IBM Tivoli Monitoring.
- Chapter 2, “Planning your OMEGAMON XE on z/VM and Linux configuration,” on page 15 describes the components of OMEGAMON XE on z/VM and Linux. Using this chapter, you will gather the information you need to make decisions about your configuration and about the deployment process.
- Chapter 3, “Installation preparation,” on page 21 describes the prerequisite software and hardware for the shared common technology components. This chapter also explains the contents of the CDs and tapes that are included in the product package.
- Chapter 4, “Upgrading IBM Tivoli Monitoring,” on page 33 describes topics to consider if you currently have one or more monitoring agents in your environment.

Part 2 - Configuring your monitoring agent

Part 2 contains the following chapters:

- Chapter 5, “Configuration required for the OMEGAMON XE on z/VM and Linux monitoring agent,” on page 37 lists the product prerequisites, and describes the configuration steps required on the systems where the monitoring agent is installed.
- Chapter 6, “Defining user IDs and security,” on page 61 provides information about user IDs and how user security is implemented in the monitoring agents.

Part 3 - Completing your configuration

Part 3 contains the following chapters:

- Chapter 7, “Performance and storage considerations,” on page 65 identifies options to consider when deploying this monitoring agent.
- Chapter 8, “Serviceability,” on page 73 describes the product features, tools, and documentation that relate to troubleshooting, problem determination, or problem source identification.

Appendixes

- “Planning worksheets” on page 79 contains each of the worksheets that you use to collect the information required to install and configure the various components.
- “Finding the information you need” on page 95 describes the kinds of tasks you perform as a user of IBM Tivoli Monitoring and of the OMEGAMON XE on z/VM and Linux agent, with pointers to the locations of the information required to perform these tasks.
- “Support for problem solving” on page 101 describes the options available for obtaining support for IBM software products.
- “Documentation library” on page 105 contains information about the publications for Tivoli OMEGAMON XE on z/VM and Linux and for IBM Tivoli Monitoring and the commonly shared components of Tivoli Management Services. This appendix also includes a section on publications and Web sites that are important for this monitoring agent.
- “Notices” on page 111 contains copyright and trademarks information.

Part 1. Planning your configuration

The chapters in this section provide the background and setup information required to install and to configure the OMEGAMON XE on z/VM and Linux monitoring agent.

- Chapter 1, “Overview of the OMEGAMON XE on z/VM and Linux monitoring agent,” on page 3 introduces the OMEGAMON XE on z/VM and Linux monitoring agent. This chapter also describes the features and components of IBM Tivoli Monitoring.
- Chapter 2, “Planning your OMEGAMON XE on z/VM and Linux configuration,” on page 15 describes the components of OMEGAMON XE on z/VM and Linux. Using this chapter, you will gather the information you need to make decisions about your configuration and about the deployment process.
- Chapter 3, “Installation preparation,” on page 21 describes the prerequisite software and hardware for the common shared technology components. This chapter also explains the contents of the CDs and tapes that are included in the product package.
- Chapter 4, “Upgrading IBM Tivoli Monitoring,” on page 33 describes topics to consider if you currently have one or more monitoring agents in your environment.

Chapter 1. Overview of the OMEGAMON XE on z/VM and Linux monitoring agent

The OMEGAMON XE on z/VM and Linux monitoring agent is a member of the IBM Tivoli Management Services family, a suite of products that monitors and manages systems and network applications on a variety of operating systems. OMEGAMON XE on z/VM and Linux gives you the ability to view data collected from multiple systems. With this monitoring agent, you can view z/VM data obtained from the Performance Toolkit for VM (also called the Performance Toolkit). You can also display views of Linux on IBM System z performance data. With this dual capability, you can solve problems quickly, and you can better manage a complex environment.

Use this monitoring agent to take advantage of the flexibility of Linux with the reliability and performance of mainframe servers. With z/VM virtualization technology, you can consolidate a significant number of Linux servers, thus minimizing costs and maximizing the availability of your mission-critical applications. Another significant benefit of virtualization is the reduction of power usage, thus contributing to a green environment and lessening your energy costs.

Predefined workspaces enable you to start monitoring your enterprise as soon as the OMEGAMON XE on z/VM and Linux software is installed and configured. The user interface supports several formats for viewing data, such as graphs, bar charts, and tables.

The displayed information enables you to:

- Collect and analyze reliable, up-to-the-minute data that allow you to make faster, better informed operating decisions.
- Identify, isolate and correct problems across z/VM and Linux instances quickly and easily.
- Manage your applications from a single point to identify problems at any time.
- Track performance across multiple platforms and systems.
- View and monitor workloads for virtual machines and workload groups, as well as analyze LPAR usage.
- View reports on z/VM usage of resources, such as CPU utilization, storage, I/O statistics, minidisk cache, and wait statistics on processor spin locks.
- View reports on Linux on System z usage of resources, such as paging data, CPU usage, and system usage by Linux guests.

With OMEGAMON XE on z/VM and Linux, you can set threshold levels to alert you when system conditions reach the defined thresholds. You can use advanced monitoring facilities that include:

- Predefined situations based on thresholds to raise different types of alerts. You can customize the predefined situations to meet the needs of your enterprise.
- Navigation to greater levels of detailed performance data. For Linux guests, this monitoring agent provides links to Tivoli Monitoring Agent for Linux OS workspaces directly from the Tivoli OMEGAMON XE on z/VM and Linux workspaces. See “Dynamic linking to cross-product workspaces” on page 9 for details.
- Expert advice on how to identify and to resolve performance problems.

With this advanced monitoring approach on the Tivoli Enterprise Portal interface, you can manage multiple z/VM systems, as well as systems on other platforms, such as Windows, UNIX[®], and z/OS[®].

From the Tivoli Enterprise Portal interface, you can also connect to the z/VM host system by means of TCP/IP to access the Performance Toolkit. Once in the z/VM Performance Toolkit, you can navigate to specific screens to view more detailed data and to investigate a problem. You use the **System Terminal** workspace to connect to the z/VM host system.

Tivoli Enterprise Portal

The OMEGAMON XE on z/VM and Linux product has a flexible, easy-to-use Java-based interface called the Tivoli Enterprise Portal, which transforms system data into the business knowledge that you can use to run your enterprise. With OMEGAMON XE on z/VM and Linux, you can also set threshold levels and flags as desired to alert you when systems reach critical points.

Figure 1 shows the Tivoli Enterprise Portal application window for OMEGAMON XE on z/VM and Linux. Tivoli Enterprise Portal presents information in a single window comprising a navigation tree and a workspace:

- The *navigation tree* in the upper left corner of Figure 1 alerts you to events using indicator icons and sound. As you move up the navigation tree hierarchy, multiple situation events are consolidated to show only the indicator with the highest severity: critical, followed by warning, and informational.
- *Workspaces* such as the one shown in Figure 1 can be divided into multiple *views* containing reports in the form of tables and charts, emulator views, Web browsers, text boxes, graphic views, and event message logs.

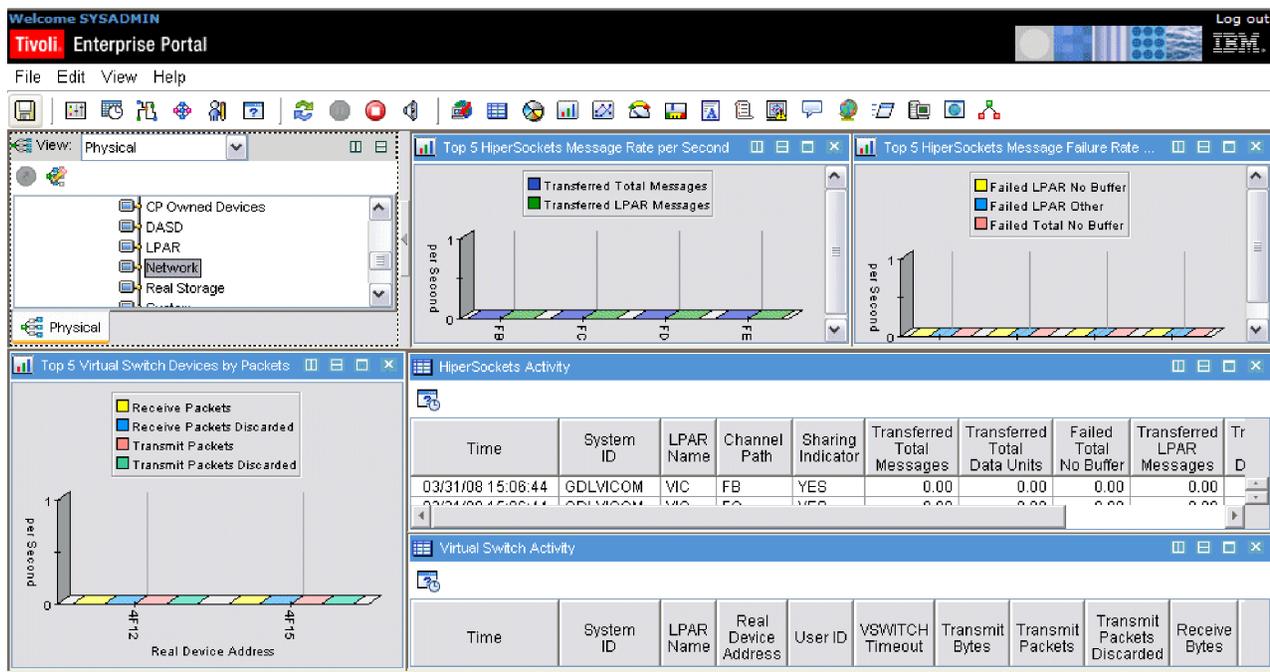


Figure 1. Tivoli Enterprise Portal sample workspace for OMEGAMON XE on z/VM and Linux

You can use the OMEGAMON XE on z/VM and Linux features to perform the following tasks:

- Monitor all systems from a single, integrated browser-based interface that you can customize with filters to display only the data you want to see
- View comprehensive online reports about system conditions
- Define your own queries, using the attributes provided by the monitoring agent, to monitor conditions and data and customize workspaces
- Create *situations* that let you set up monitoring for particular conditions and flag the condition with an alert when detected

Note: If you create a situation and specify that an Action be run, be aware of the following:

- For this monitoring agent, select **Run the action at the Managed System (Agent)** for the **Where should the action be run (Performed)** option.
 - When you select **System Command** on the Action tab, this field displays for you to type a command to issue at the system. Be sure to prefix all OMEGAMON XE on z/VM and Linux commands with **VL:**.
- Trace the causes leading up to an alert
 - Send an alert when an event has occurred on a managed system
 - Customize workspaces to meet the unique needs of your enterprise
 - View performance in real-time and usage trends from a historical perspective
 - Establish performance thresholds, which highlight displayed values with color indicators when the value falls outside the threshold limits
 - Create and send commands to systems in your managed enterprise by means of the *Take Action* feature. The Take Action feature lets you enter a command or select from one of the predefined commands for your product and run it on any system in your managed network either manually or automatically in response to reported conditions. This monitoring agent has specific requirements for enabling Take Action commands. See “Step 11. Enabling Take Action commands (optional)” on page 53.
 - Embed information about potential fixes in the product interface using *Expert Advice*, which can be edited to include knowledge and solutions specific to your environment

What's new in Version 4.2.0

The following changes have been made to this Tivoli OMEGAMON XE on z/VM and Linux V4.2.0 release:

- Replacement of the z/VM Systems default workspace with the z/VM Systems System Health workspace.
- Overall system view.

The sections that follow provide more detail about these changes.

Replacement of the z/VM Systems: default workspace with the z/VM Systems: System Health workspace

The **z/VM Systems** branch is the top level of the Navigation tree for the z/VM operating system. When you click **z/VM Systems**, the System Health workspace for this level of the tree is displayed. This workspace provides data for all of the z/VM systems that are registered with the Tivoli Enterprise Monitoring Server. Use this workspace to view the overall health of all your z/VM systems.

Overall system view

A new attribute group called KVLSystem2 attributes has been added. This attribute group provides system-level data on CPU utilization by the z/VM Control Program and by the z/VM virtual machines, on free-storage management, on paging rates, and on user activity. New attributes have also been added to the following attribute groups: KVLSystem, KVLLPAR Info, KVLSMinidisk Cache, and KVLUser Workload. For more information on the attributes for the OMEGAMON XE on z/VM and Linux product, see the IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide.

New product-provided situations

The following table contains the new situations provided by this version of the monitoring agent:

Table 1. New predefined situations

Navigator item	Name of situation	Column name and initial conditional value	State	Runs at startup (Yes or No)
z/VM Linux Systems	ZVM_Storage_Overcommit_High	KVLSystem2.Real_Storage_Overcommit > 3.0 and KVLSystem2.Real_Storage_Overcommit <= 4.0	High (Warning)	No
z/VM Linux Systems	ZVM_Storage_Overcommit_Critical	KVLSystem2.Real_Storage_Overcommit > 4.0	Critical	No
System	ZVM_Eligible_List_High	KVLSystem.Eligible_Users > 5	High (Warning)	No

Note: If you customize any product-provided situation, you may lose any changes you make when these situations are modified by future application of maintenance. To retain the changes you make to the existing situations, make a copy of the situation using a unique name, and modify the copy instead of the original. That way your situations will not be overwritten by maintenance.

Tivoli Management Services components

All products that display data in the Tivoli Enterprise Portal use a common shared technology. The common shared technology is known as Tivoli Management Services. Tivoli Management Services consists of Tivoli Enterprise Portal, Tivoli Enterprise Portal Server, Tivoli Enterprise Monitoring Server and Tivoli Data Warehouse.

IBM Tivoli Monitoring V6.x is part of Tivoli Management Services. The IBM Tivoli Monitoring product media includes the Tivoli Management Services components and the IBM Tivoli Monitoring agents that monitor distributed systems and applications. When you purchase IBM Tivoli Monitoring, you are licensed to use the IBM Tivoli Monitoring components and the distributed monitoring agents. When you purchase an OMEGAMON XE V4.x product, the IBM Tivoli Monitoring V6.x distribution media are included in the product package. The required versions of the common shared technology components are included on this media. This media should be used to update the common shared technology components if the required level has not already been installed.

For information about the IBM Tivoli Monitoring components, see “Understanding and designing your configuration” on page 16. For information about supported operating systems for each component, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

Important: All Tivoli products that utilize the IBM Tivoli Monitoring V6 components are impacted by Daylight Savings Time (DST). Several countries have announced changes to their Daylight Savings Time (DST) policies and/or changes to their Time Zones. In order to assist customers with understanding the impact of these changes to their Tivoli products, a Knowledge Collection with associated links has been created.

For more information, refer to *Knowledge Collection: Non-US Daylight Savings Time (DST) and Time Zone changes impact on Tivoli Products*, found within the **Technotes** link at the following URL address:

<http://www-306.ibm.com/software/sysmgmt/products/support/IBMTivoliMonitoringV6.html>

IBM Tivoli Monitoring features

The following features are available with IBM Tivoli Monitoring and the Tivoli Enterprise Portal:

- **Customized workspaces for each information group:** Tivoli Enterprise Portal retrieves data from the monitoring agent and displays the results in the workspace in the form of charts and tables. You can start monitoring activity and system status immediately with the predefined workspaces. With just a few clicks of the mouse, you can tailor your own workspaces to look at specific conditions, display critical threshold values in red, and filter incoming data according to your needs.
- **Workspace views:** Each workspace consists of one or more views. There are several types of views:
 - *Table views* display data in table format where rows represent monitored resources and columns represent data collected for each resource.
 - *Chart views* display data in graphical formats. Pie, bar, and plot charts and a gauge format are supported.
 - *Take action view.* With this view, you can enter a command or select a predefined command, and run it on any system in your managed network.
 - *Message log view* shows the status of the situations running on your managed network.
 - *Notepad view* opens a simple text editor for writing text that can be saved with the workspace.
 - *Terminal view* starts an emulator session for working with z/OS or z/VM applications.
 - *Browser view* opens the integrated Web browser.
- **Navigator views or navigators** provide hierarchical views of the systems, resources, and applications you are monitoring. Navigators help you structure your enterprise information to reflect the interests and responsibilities of the user. OMEGAMON XE comes with a default navigator called the physical navigator.

OMEGAMON Dashboard Edition (DE) on z/OS comes with the same default navigator. However, you can create additional navigators that display enterprise information representing your business systems. OMEGAMON DE on z/OS is separately orderable.

Note: You cannot use OMEGAMON DE on z/OS with OMEGAMON XE on z/VM and Linux, unless you are licensed to do so. However, OMEGAMON DE for distributed products is included as part of the OMEGAMON XE on z/VM and Linux monitoring agent installation.

- **Linked workspaces:** If you often go from one workspace to another, you can build a link between them to speed the transition. You can also build links that originate from a table or from a bar or pie chart, and use relevant data from the source table or graph to determine the target workspace.
- **Historical data collection:** Historical data collection provides the ability to store data collected over time. Without historical data collection, all data presented on the Tivoli Enterprise Portal is from the most recent data collection. Collecting historical data makes it possible for you to view graphs of data over time. Data collected for previous time periods is available only if historical data collection is enabled.

You configure historical data collection by invoking the **Historical Collection Configuration** panel from the Tivoli Enterprise Portal. This panel asks you to specify the collection of historical data from either the Tivoli Enterprise Monitoring Server or from the monitoring agent. You also select which tables and columns are to be collected and the interval at which to collect the historical data.

Configuration of historical data collection at the Tivoli Enterprise Portal allows you to specify the following options:

- Attribute groups to be collected
- Historical data collection interval
- Storage location for short term history. Short term historical data can be stored either on the management agent or on the monitoring server. Storing it on the management agent is recommended. Display of short-term history is limited to 24 hours.
- Long-term historical data warehousing interval (if you choose to write your data to the Tivoli Data Warehouse)
- Data summarization and pruning are available if the Tivoli Data Warehouse is enabled.

With Tivoli Data Warehouse, you can analyze historical trends from monitoring agents. The Tivoli Data Warehouse uses a DB2®, Oracle, or Microsoft SQL Server database to store historical data collected across your environment. You can generate warehouse reports for short term and long term data through the Tivoli Enterprise Portal. You can also use third-party warehouse reporting software, such as Crystal Reports or Brio, to generate long term data reports. Warehouse reports provide information about the availability and performance of your monitoring environment over a period of time.

The Tivoli Data Warehouse uses the Warehouse Proxy agent to move data from monitoring agents or the monitoring server to the data warehouse database. The Warehouse Proxy is an ODBC export server for warehousing historical data. It is a special agent that uses an ODBC connection to transfer historical data collected from agents to a database. You can then analyze this data using the workspaces in the Tivoli Enterprise Portal or any third-party software. For information on installing the Warehouse Proxy agent, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

The Warehouse Summarization and Pruning agent provides the ability to customize the length of time for which to save data (pruning) and how often to compress data (summarization). For information about the Warehouse Summarization and Pruning agent, see the *IBM Tivoli Monitoring Administrator's Guide*. If you do not intend to use historical reporting or save historical data to a database for reference, then you do not need to install or configure the Warehouse Proxy. For information on installing the Tivoli Data Warehouse, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

- **Custom queries:** Every monitoring agent comes with a set of predefined queries. These queries tell the monitoring server what monitoring data to retrieve from the agent for the chart or table view. You can create your own queries to specify exactly which attributes to retrieve, thus saving valuable resources. For example, you can build a filter into the LPAR query to retrieve only records from a particular logical partition. Additionally, you can write SQL queries to ODBC data sources and display the results in any chart or table. This enables you to show monitoring data and data from other sources (such as third-party databases) in a single workspace.
- **Interaction with systems from your console:** Use the Take Action feature to enter a command or to select a predefined command, and run it on any system in your managed network.
- **Monitor system conditions and send alerts:** You can use the situation editor to create situations. A situation notifies you when an event occurs on a managed system. The monitoring server sends an alert when the conditions in a situation are evaluated to be true. The alert is displayed on the Tivoli Enterprise Portal with visual and sound indicators. An event is only generated the first time a situation becomes true, and then it is reset. This monitoring agent comes with a set of predefined situations.

Note: When you create or modify situations for this monitoring agent and you choose to issue commands on z/VM when the situation becomes true, note the following items:

- For the **Where should the action be run (Performed)** option, select **Run the action at the Managed System (Agent)**.
- When you select **System Command** on this tab, this field displays for you to type a command to issue at the z/VM system. Be sure to prefix all OMEGAMON XE on z/VM and Linux commands with **VL:**.
- **Managed system lists:** You can create and maintain named lists of managed systems that can be applied to:
 - Situation distribution lists
 - Policies correlated by business application group
 - Queries
 - Customer Navigator-managed system assignments
- **User administration:** Tivoli Enterprise Portal provides the user administration feature for adding new user IDs, complete with selectable permissions for the major features and specific managed systems.
- **Universal Agent support:** The Universal Agent is an agent you can configure to monitor any data you collect. Use this agent to integrate data from virtually any operating system and any source, such as custom applications, databases, systems, and subsystems. Your defined data providers are listed in the Navigator, and default workspaces are automatically created for them.

Dynamic linking to cross-product workspaces

With dynamic workspace linking, you can easily navigate between workspaces that are provided by multiple products. This feature aids problem determination and improves integration across the monitoring products. You can then more quickly determine the root cause of a problem. Using the predefined cross-product links provided by the OMEGAMON XE products, you can obtain additional information about systems, subsystems, resources, or network components that are being monitored by other monitoring agents. When you right-click on a link, the list of links is displayed. This list might contain links to workspaces provided by one or more monitoring products. The product you are linking to must be installed and configured.

Note: Your Tivoli Enterprise Portal user ID must be authorized to access the target product. Otherwise links to workspaces in the targeted product are not included in the list. Choose a workspace from the list to navigate to that workspace. By linking to the target workspace in context, you receive additional information that is related to the system, subsystem, or resource you are currently viewing. If you choose a workspace from the list and the target workspace is not available, you receive message KFWITM081E.

This monitoring agent contains dynamic workspace links to workspaces in the Tivoli Monitoring Agent for Linux OS. Refer to the workspace descriptions in the *IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide* for information about the predefined links provided with each workspace. See the *IBM Tivoli Monitoring: Agent for Linux OS User's Guide* for information about the workspaces available with this monitoring agent. If you have the Tivoli Monitoring Agent for Linux OS installed and you would like to link to workspaces in that monitoring agent, you must configure an environment variable in the Tivoli Monitoring Agent for Linux OS. See “Step 10. Enabling dynamic workspace linking” on page 52 for details.

OMEGAMON XE on z/VM and Linux monitoring agent architecture overview

Monitoring agents monitor and collect performance data from managed systems. In general, monitoring agents monitor systems, subsystems, resources, and applications on the system where they are installed. Monitoring agents provide data and performance information to the monitoring server and receive instructions from the monitoring server. You can also use the monitoring agents to issue commands to the system or application you are monitoring.

The OMEGAMON XE on z/VM and Linux monitoring agent is installed and runs as a process on a Linux on zSeries® guest. The agent obtains z/VM data and Linux on System z data from the Performance Toolkit running on z/VM. The acquired data displays in the Tivoli Enterprise Portal. This capability allows for quick problem resolution and enables support personnel to more easily work across platform boundaries. See “Software and hardware prerequisites” on page 21 for information on the software and hardware prerequisites.

Figure 2 on page 10 depicts a sample OMEGAMON XE on z/VM and Linux environment. In this sample environment, one z/VM image is running with three primary guests. The primary guests include two Linux on zSeries guests and one guest running z/OS. The monitoring server, the portal server, the Tivoli Enterprise Portal, and the data warehouse are running on a separate host.

In this sample environment, one of the Linux systems is running the OMEGAMON XE on z/VM and Linux monitoring agent. Additionally, both Linux systems are running the IBM Tivoli Monitoring Agent for Linux OS. All of the monitoring agents are transferring requests and data using the shared technology components.

The environment includes a Conversational Monitor System (CMS) guest that is running the Performance Toolkit. The data are collected from the z/VM operating system. The data are written to a discontinuous saved segment (DCSS) on z/VM. The OMEGAMON XE on z/VM and Linux monitoring agent reads the data in the DCSS using the z/VM DCSS device driver support in Linux for zSeries, and displays the data in the Tivoli Enterprise Portal. This monitoring agent provides a predefined default DCSS called **PERFOUT**

that can be used in most environments. See Chapter 5, “Configuration required for the OMEGAMON XE on z/VM and Linux monitoring agent,” on page 37 for details.

Lastly, another CMS guest is used as a VM command server. The VM command server receives command requests from the Tivoli Enterprise Portal by means of the monitoring agent. Use the Take Action feature to execute VM and/or CMS commands to the z/VM operating system. See “Step 11. Enabling Take Action commands (optional)” on page 53.

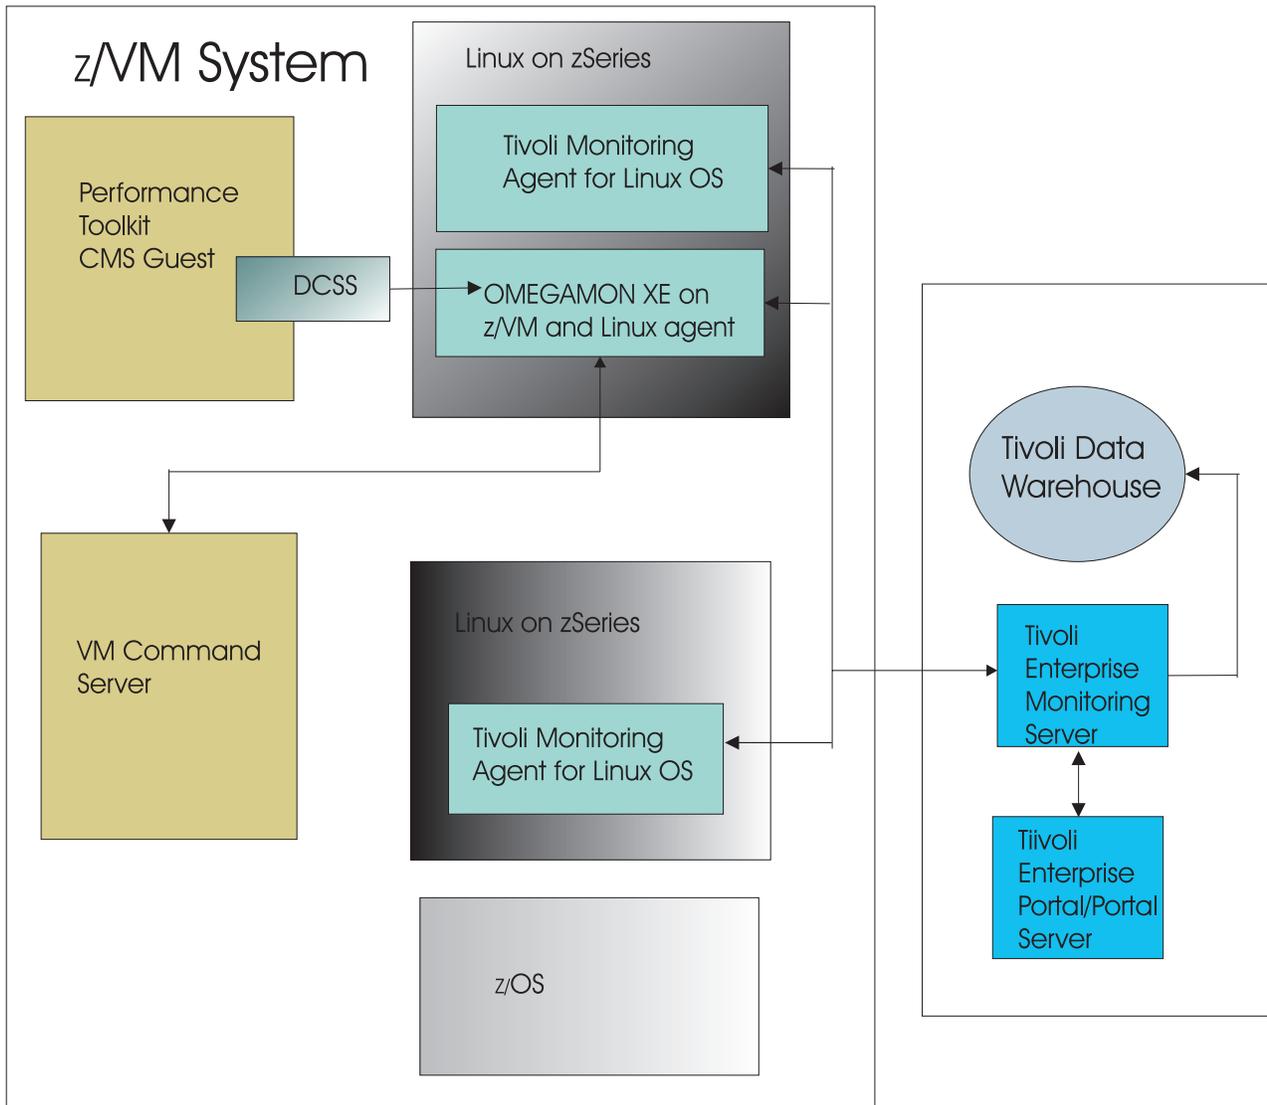


Figure 2. OMEGAMON XE on z/VM and Linux architecture

IBM Tivoli OMEGAMON XE zSeries products

The IBM Tivoli OMEGAMON XE zSeries suite of products includes solutions for z/OS-based applications and Linux on System z applications, database products, applications such as CICS®, storage, and networks. For additional information on any of these agents, refer to the product documentation.

The IBM Tivoli OMEGAMON XE zSeries management solutions help you achieve a true on-demand computing environment. These products can help you meet the demands of increasing data center volume, complexity, and volatility by giving your IT team the tools to quickly identify, isolate, and fix problems before they impact customers.

The following table lists and describes some of the IBM Tivoli OMEGAMON XE zSeries agents. You can see the complete list of IBM Tivoli OMEGAMON XE zSeries agents at <http://www.ibm.com/software/tivoli/solutions/availability/>.

Table 2. IBM Tivoli OMEGAMON XE zSeries products

Product	Description
Tivoli OMEGAMON XE for CICS	Helps you proactively manage complex CICS systems (including CICS in a Parallel Sysplex environment) to achieve high performance and avoid downtime. With this interface, you can clearly see and understand application and system events. Designed to enable you to detect problems quickly and take action in real time to speed problem resolution.
Tivoli OMEGAMON XE for DB2 Performance Monitor/Expert on z/OS	Helps you proactively manage your DB2® mainframe environment and tune for optimal performance. These products bring the strength of both OMEGAMON for DB2 and DB2 Performance Monitor/Expert into a new set of DB2 monitors. The difference between the Monitor and the Expert product offerings is that the Expert provides additional expert analysis functions like buffer pool analysis and simulation as well as specific expert rule-of-thumb and SQL performance queries run against data stored in the performance warehouse.
Tivoli OMEGAMON XE for IMS™ on z/OS	Helps you optimize the performance and availability of your vital IMS systems. Provides a single point of control over IMS in parallel Sysplex environments and reports on performance of coupling facility structure statistics, shared queue counts, database lock conflicts and a number of other key IMS attributes that help you stay ahead of potential delays or outages. The Tivoli OMEGAMON XE for IMS product combines the functions offered by OMEGAMON XE for IMS and OMEGAMON XE for IMSplex into a single product. This agent includes the following additional features: <ul style="list-style-type: none"> • Lock table enhancements for additional owner and identification information • Extended TRF and TRF Extractor functionality and reporting • OTMA extensions and COLD queue reporting • Shared Queues extensions • IMS Connect CPU time statistics
Tivoli OMEGAMON XE for Mainframe Networks	Monitors the TCP/IP and SNA resources on a z/OS system. Collects network performance data from a z/OS system and presents the information through the Tivoli Enterprise Portal. Raises alerts within the user interface, and can export them to event-receiving products such as Tivoli Event Console and NetView® for z/OS. This agent includes the following additional features: <ul style="list-style-type: none"> • Collects data through the z/OS Communications Server Network Management Interfaces (NMI) plus continuation of some data collection through SNMP for more efficient collection. • Provides IP performance data equivalent to the IBM Tivoli Monitoring for Network Performance V2.1 product and a superset of SNA performance data. • Provides SNA performance data to sustain operations until migration to TCP/IP. • Provides initial integration and interoperability with NetView for zSeries V5.2. • Enables or disables collection of categories of data through product configuration and dynamically through the z/OS MODIFY command.
Tivoli OMEGAMON XE for Storage on z/OS	Provides a comprehensive monitor for z/OS I/O subsystem performance and storage availability. Designed to manage the performance and availability of mainframe-attached storage, including hard disk storage and tape devices and the data sets that reside on them. The product also features in-depth analysis of two important IBM storage software components: <ul style="list-style-type: none"> • The Data Facility Systems Managed Storage, which manages the service levels and priorities of data sets based on user created storage goals. • The Data Facility Hierarchical Storage Manager, which manages backup of data based on usage patterns. <p>Additional function includes exploitation of the new IBM DS6000™ and DS8000® storage devices and dataset masking functionality for dataset mask group capabilities and ease of use.</p>

Table 2. IBM Tivoli OMEGAMON XE zSeries products (continued)

Product	Description
Tivoli OMEGAMON XE on z/OS	<p>Combines the monitoring capabilities of OMEGAMON XE for Sysplex, OMEGAMON XE for OS/390®, and OMEGAMON XE for IBM Cryptographic Coprocessors. This product enables you to effectively monitor the availability, performance, and resource utilization of sysplexes and the individual z/OS systems that participate in them. Provides comprehensive performance information covering Sysplex level components such as Workload Manager, Coupling Facility, Cross System Coupling Facility (XCF), Global Enqueue, and shared DASD as well as detailed system-level information. This agent includes the following additional features:</p> <ul style="list-style-type: none"> • Migration of key features from the previous OMEGAMON II® for MVS™ product into the OMEGAMON XE on z/OS product, including detailed CSA usage by address space and Inspect functionality. • Address space level CPU usage times and percentages • Enhanced system CPU utilization reporting • Enhanced zSeries Application Assist Processor (zAAP) processor usage and reporting

Additionally, IBM Tivoli OMEGAMON Dashboard Edition (DE) on z/OS is a package of components that provides an integrated view of your mainframe enterprise and the power to take corrective action when problems threaten system and application availability. This interface also gives you workflow policies to define and run complex automation scenarios. With OMEGAMON DE, you can combine data from different agents into a single OMEGAMON DE screen for data integration. The components in the package include OMEGAVIEW and OMEGAVIEW II® for the Enterprise. If you are an OMEGAVIEW customer, you can continue to use these products until you are ready to begin using the most current, improved OMEGAMON DE features. IBM Tivoli OMEGAMON DE on z/OS is separately orderable.

Note: You cannot use OMEGAMON DE on z/OS with OMEGAMON XE on z/VM and Linux, unless you are licensed to do so. However, OMEGAMON DE for distributed products is included as part of the OMEGAMON XE on z/VM and Linux monitoring agent installation.

Serviceability

Serviceability is defined as product features, tools, and documentation relating to troubleshooting, problem determination, or problem source identification. On a broader level, it can also encompass service offerings or processes that make the product more serviceable. To make the OMEGAMON XE zSeries products serviceable, log files containing messages and trace information are provided in a common fashion across all the monitoring agents and the components of the common shared technology.

The log files are available to you to help in resolving problems encountered while using the products. IBM Software Support may request some or all of these files while investigating a problem you have reported. For more information about service issues relating to the components of Tivoli Management Services (such as Tivoli Enterprise Monitoring Server and Tivoli Enterprise Portal), refer to *IBM Tivoli Monitoring Troubleshooting Guide*. For problem determination relating to this monitoring agent, see the *IBM Tivoli OMEGAMON XE on z/VM and Linux Troubleshooting Guide*. See also Chapter 8, “Serviceability,” on page 73.

Standards supported

Tivoli OMEGAMON XE zSeries products provide a number of integration facilities and adhere to a range of industry standards to make integration with other applications easier for you. These products use industry-standard languages and protocols to facilitate integration with third-party components and tools. The product also uses strategic IBM and Tivoli tools and platforms.

These standards and platforms include the following:

- A Web-based user interface implemented with industry-standard Web content languages, such as Java™, XML, and HTML.

- OMEGAMON Web Services, an open standards-based interface to IBM Tivoli Monitoring using Simple Object Access Protocol (SOAP) requests. Any OMEGAMON XE monitoring agent can be dynamically queried, so that performance and availability can be processed by other applications.
- Simple Network Management Protocol.
- Web Services and Web Management Interface (WMI) standard.
- TCP/IP-based communications between components and systems.
- Support for the DB2 product, an industry-standard relational database.
- Use of Structured Query Language (SQL '92, ISO/IEC 9075:1992), the standard interface for relational database access.
- Use of standard shell scripts, SMP/E, and VMSES/E to assist in installation.
- Installation on a Windows system using the industry-standard InstallShield.
- Installation on a UNIX system using a proprietary Java-based installation tool.

Interoperability with other products

Interoperability is the capability of an application to integrate with other IBM and non-IBM applications that are used in the same customer environment. Tivoli OMEGAMON XE zSeries products are compatible with each other and can coexist in a single OMEGAMON XE environment (that is, with a common Tivoli Enterprise Monitoring Server). These products also interoperate with Tivoli Enterprise Monitoring Agents running on distributed systems and communicating through the same monitoring server.

Chapter 2. Planning your OMEGAMON XE on z/VM and Linux configuration

In this chapter, you will learn about the components of OMEGAMON XE on z/VM and Linux, and gather the information you need to make decisions about your configuration. This chapter is especially useful if you are new to OMEGAMON XE zSeries monitoring agents.

Before you begin the tasks of configuring OMEGAMON XE on z/VM and Linux, be sure to complete these prerequisite steps:

1. Read the *IBM Tivoli OMEGAMON XE on z/VM and Linux Program Directory* for this monitoring agent and complete all of the installation requirements listed there.
2. Read the *IBM Tivoli Monitoring Installation and Setup Guide* and complete all of the installation requirements listed there.
3. Read this chapter to determine how you want your configuration to look. For example, you must decide where you want to deploy Tivoli Enterprise Monitoring Servers and monitoring agents.

To prepare for the installation and configuration process, fill out all of the following worksheets:

- “Worksheet: Your overall configuration” on page 79
- “Worksheet: Your monitoring agent configuration” on page 80
- “Worksheet: Planning communication protocols for the monitoring agent when the monitoring server is on a distributed system” on page 81
- “Worksheets: Information to gather when configuring your portal server on Windows or Linux” on page 82
- “Worksheets: Information to gather when configuring your monitoring server on a distributed system” on page 84
- “Worksheets: Information to gather when putting your hub monitoring server on a z/OS system” on page 86
- “Worksheets: Information to gather when configuring your portal desktop client on Windows or on Linux” on page 91
- “Specifying communication protocols between components” on page 93

Understanding and designing your configuration

IBM Tivoli Monitoring is considered a client-server-agent implementation.

Figure 3 shows the major components of IBM Tivoli Monitoring.

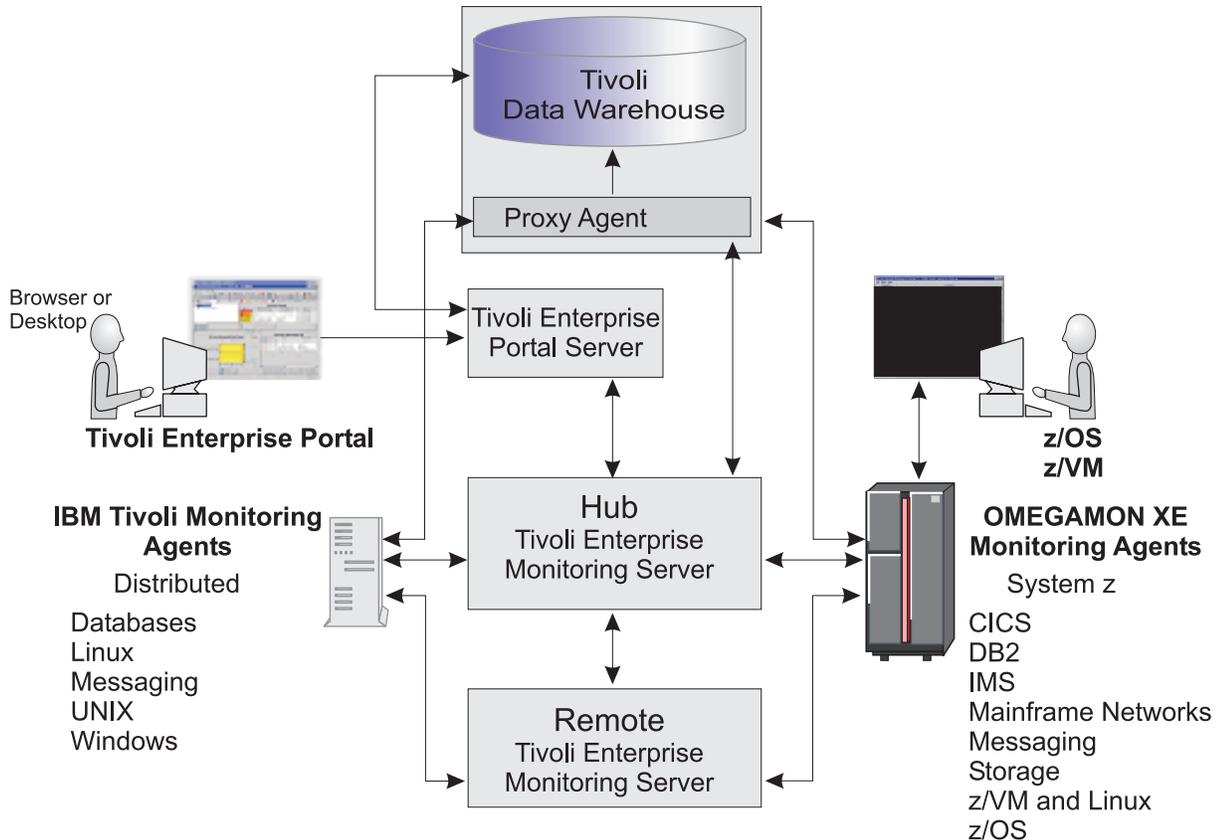


Figure 3. IBM Tivoli Monitoring components

The components include:

- “Tivoli Enterprise Portal and Tivoli Enterprise Portal Server”
- “Tivoli Enterprise Monitoring Servers - hub and remote” on page 17
- “OMEGAMON XE on z/VM and Linux monitoring agent architecture overview” on page 9
- “Tivoli Enterprise Console” on page 19
- “Tivoli Data Warehouse” on page 19

Tivoli Enterprise Portal and Tivoli Enterprise Portal Server

The *Tivoli Enterprise Portal client* (portal client) is the user interface for the IBM Tivoli Monitoring products. The portal client is a thin Java application that communicates with the Tivoli Enterprise Portal Server to send requests and to retrieve data. You can access all portal client function through the desktop client or through an Internet Explorer browser connected to an embedded Web server in the Tivoli Enterprise Portal Server.

- The desktop client requires that the Tivoli Enterprise Portal component be installed and maintained on each user’s desktop. The desktop client can run on Windows or Linux (Intel® Linux only - Red Hat or SUSE).

- Using the browser client, you can leverage an existing deployment of Internet Explorer and Java Runtime Environment without installing the client component on every user's workstation. The browser client can run on Windows only, with Internet Explorer 6 and Mozilla Firefox supported.

If you use the use the Tivoli Enterprise Portal in browser mode and you migrated to a new release or a fix pack of this monitoring agent, you may encounter problems seeing data in the new workspaces. This is because the Java Plug-in cache for the browser client needs to be cleared between releases. Clearing the Java Plug-in cache removes old versions of the Tivoli Enterprise Portal JAR files and prevents exception messages. See the *IBM Tivoli Monitoring Installation and Setup Guide* for information on installing and using the browser client, and for instructions on clearing the Java Plug-in cache.

See the *IBM Tivoli Monitoring Installation and Setup Guide* for complete information about operating system version support.

The *Tivoli Enterprise Portal Server* (portal server) is a Java application server that enables retrieval, manipulation, and analysis of data from agents. The portal server connects to both the portal client and the Tivoli Enterprise Monitoring Server. The portal server holds all the information required to format the workspaces viewed in the portal client.

The portal server communicates with the clients (default port is 1920) and with the hub Tivoli Enterprise Monitoring Server (default port for IP network is 1918). You can provide fault tolerance by connecting more than one portal server to the same hub Tivoli Enterprise Monitoring Server.

Decision point:

How do you choose between Windows and Linux for installation of the portal server and portal desktop client?

You must base the decision about whether to use Linux or Windows for the portal server and desktop client entirely on conditions and preferences at your site, such as:

- The operating systems already in use in the existing environment
- Familiarity and comfort level with the Windows and Linux operating systems
- Whether you want to bring additional operating systems into your site's current configuration

Note that you can run with mixed Windows and Linux portal server and portal client components. For example, you can have a desktop client on Linux and a portal server on Windows, or a desktop client on Windows and a portal server on Linux.

See the *IBM Tivoli Monitoring Installation and Setup Guide* and the *IBM Tivoli Monitoring Administrator's Guide* for information about planning for IBM Tivoli Monitoring components on Windows or Linux.

Tivoli Enterprise Monitoring Servers - hub and remote

All requests and data for OMEGAMON XE monitoring agents, such as OMEGAMON XE on z/VM and Linux, flow through a hub Tivoli Enterprise Monitoring Server (monitoring server). The monitoring server component performs the following tasks:

- Retrieves data from the monitoring agents and delivers data to the portal server.
- Sends alerts to the portal server when conditions specified in situations are met.
- Receives commands from the portal client and passes them to the appropriate monitoring agents.
- Optionally, provides a repository for short-term historical data.

You can install this component on z/OS, Windows, Linux Intel, Linux on zSeries, and some UNIX operating systems. See the *IBM Tivoli Monitoring Installation and Setup Guide* for a complete list of supported platforms.

Decision point:

Must you install a monitoring server on z/OS, Windows, Linux on zSeries, or UNIX systems?

Many organizations prefer the reliability and availability characteristics of the z/OS platform for the monitoring server. However, if you have other OMEGAMON XE monitoring agents on your Windows or UNIX systems, you might prefer Windows or UNIX platforms.

The two basic types of monitoring servers are *hub* and *remote*:

- The *hub* monitoring server is the focal point for managing your environment. You can configure only one hub monitoring server, which communicates with the portal server, with monitoring agents, and optionally with monitoring servers running remotely.
- You can optionally configure a *remote* monitoring server to distribute the workload of the hub monitoring server, but it is not required.

Each remote monitoring server must be installed on its own computer or workstation. A remote monitoring server communicates with the hub monitoring server and with monitoring agents running on the same or different systems. Note that a remote monitoring server is remote only with respect to the hub monitoring server, not necessarily with respect to the monitoring agents. A monitoring agent can be installed on the same system as a remote monitoring server. The monitoring server is then local to the monitoring agent, but it is still a remote monitoring server. For instructions on configuring remote monitoring servers, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

Decision point:

Must you configure a remote monitoring server in your environment?

A remote monitoring server is designed to offload work from the hub. Whether or not your hub becomes overloaded enough to slow down hub processing of situations and other data depends on the complexity of your environment. The following factors tend to boost strain on the hub and increase the likelihood that you might want a remote server to help out the hub:

- You are monitoring many Linux systems and z/VM systems. The more monitoring agents you have installed on Linux on zSeries, the more work there is for the hub.
- You are monitoring many situations. OMEGAMON XE on z/VM and Linux does not come with a great many situations to consume hub cycles, so unless you have other monitoring agents with lots of situations, this is probably not a factor that will push you into needing remote monitoring servers.

Configuring a remote monitoring server can also give you scalability potential and failover protection, which might be especially important when you add OMEGAMON XE on z/VM and Linux to an environment with multiple Tivoli Management Services products and agents.

For more information on these issues, see the *Deployment Guide Series: IBM Tivoli Monitoring 6.1 Redbooks*[®], accessed from the following Web site:

<http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/sg247188.html?Open>

Look for the following topics:

- Small to medium installation
- Scalability

Tivoli Enterprise Console

The Tivoli Enterprise Console synchronizes the status of situation events that are forwarded to the event server. When the status of an event is updated because of Tivoli Enterprise Console® rules or operator actions, the update is sent to the monitoring server, and the updated status is reflected in both the Situation Event Console and the Tivoli Enterprise Console event viewer.

Adding the Tivoli Enterprise Console product, which is designed specifically for enterprise computing environments, to your IBM Tivoli Monitoring environment adds powerful event management capabilities to IBM Tivoli Monitoring.

A .baroc file is shipped with this product. The .baroc file contains the Tivoli Enterprise Console classes that are generated by IBM Tivoli Monitoring. In order to view the event data in the Tivoli Enterprise Console, you need to install this .baroc file on the event server. You add the .baroc file after you have added application support for the agent to the monitoring server. See the *IBM Tivoli Monitoring Installation and Setup Guide* for instructions on installing the monitoring agent .baroc files to the event server.

Tivoli Data Warehouse

The *Tivoli Data Warehouse* is a long-term data store for the performance and analysis data collected by the Tivoli monitoring agents or collected by the monitoring server. With Tivoli Data Warehouse, you can analyze historical trends from the monitoring agents. The Tivoli Data Warehouse uses a DB2, Oracle, or Microsoft SQL Server database to store historical data collected across your environment. You can generate warehouse reports for historical data through the Tivoli Enterprise Portal. You can also use third-party warehouse reporting software, such as Crystal Reports or Brio, to generate long-term data reports. Warehouse reports provide information about the availability and performance of your monitoring environment over a period of time.

The Tivoli Data Warehouse uses the *Warehouse Proxy agent* to move data from the monitoring server or from monitoring agents to the data warehouse database. The Warehouse Proxy is an ODBC export server for warehousing historical data. It is a special agent that uses an ODBC connection to transfer historical data collected from the agents or from the monitoring server to a database. You can then analyze this data using the workspaces in the Tivoli Enterprise Portal or any third-party software.

The Warehouse Summarization and Pruning agent provides the ability to customize the length of time for which to save data (pruning) and how often to compress data (summarization). As the size of the data grows, so does the need for more powerful tools to administer and manage the information. The Summarization and Pruning agent supports configurable data summarization and data pruning. With summarized data, the performance of queries can be improved dramatically. In addition, with data summarization and data pruning working together, the amount of disk space utilized can be better managed. For information about the Warehouse Summarization and Pruning agent, see the *IBM Tivoli Monitoring Administrator's Guide*.

If you do not intend to use historical reporting or save historical data to a database for reference, then you do not need to install or configure the Warehouse Proxy.

Warehouse Proxy planning

Consider the following when planning the deployment of your Warehouse Proxy agent:

- You can use only one Warehouse Proxy per hub monitoring server. However, if you have multiple hub monitoring servers, the Warehouse Proxy agents for those monitoring servers can share a single Tivoli Data Warehouse database, enabling you to consolidate your historical data in one location. To configure multiple Warehouse Proxy agents to use the same warehouse, when you create the ODBC connection for each Warehouse Proxy, point to the same remote database and connect with the same user.
- If you have a large monitoring environment, install the Warehouse Proxy and the Warehouse Summarization and Pruning agent on different computers. If the Tivoli Data Warehouse database is on a

UNIX or Linux computer, install the Warehouse Summarization and Pruning agent on the same computer as the Tivoli Data Warehouse database. However, the Warehouse Summarization and Pruning agent does not support all hardware configurations, so ensure that the computer where you install this agent meets the requirements for the agent.

- You can configure historical data collection to store data at any combination of the monitoring server or the agents. To ensure that history data are received from all sources, you must configure a common communication protocol between the Warehouse Proxy and the component that is sending history data to it (either from a monitoring server or from an agent).

For example, you might have a monitoring server configured to use both IP.UDP and IP.PIPE. In addition, one agent might be configured with IP.UDP and a second agent with IP.PIPE. In this example, configure the Warehouse Proxy to use both IP.UDP and IP.PIPE.

- The default user ID that the Warehouse Proxy uses to access the database is **ITMUser**.

For information on installing and configuring the Tivoli Data Warehouse and the Warehouse Proxy agent, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

Chapter 3. Installation preparation

This chapter includes a description of the prerequisite software and hardware for the common shared technology components of IBM Tivoli Monitoring and for OMEGAMON XE on z/VM and Linux. This chapter also explains the contents of the CDs and tapes that are included in the product package.

See the *IBM Tivoli Monitoring Installation and Setup Guide*, for instructions on installing this monitoring agent. That guide also provides instructions for installing and enabling application support for the monitoring agent on the shared technology components.

Software and hardware prerequisites

The prerequisites for the Command Processor component of this monitoring agent are described in the *IBM Tivoli OMEGAMON XE on z/VM and Linux Program Directory*. The prerequisites for installing IBM Tivoli Monitoring and the monitoring agent are detailed in the *IBM Tivoli Monitoring Installation and Setup Guide*. Details about these prerequisites, as well as additional prerequisites for this monitoring agent, are found in the sections that follow.

Required software

The software products and levels specified below are software requirements for running IBM Tivoli Monitoring and the OMEGAMON XE on z/VM and Linux monitoring agent.

For the latest information about the supported operating systems for the different IBM Tivoli Monitoring components, see the following Web address:

http://www-306.ibm.com/software/sysmgmt/products/support/Tivoli_Supported_Platforms.html

Software requirements for Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal, and Tivoli Enterprise Portal Server on distributed operating systems

Detailed software requirements for all distributed components of IBM Tivoli Monitoring (Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal, and Tivoli Enterprise Portal Server running on any of the supported UNIX operating systems or Windows versions) are found in *IBM Tivoli Monitoring Installation and Setup Guide*.

For information on installing the Tivoli Enterprise Monitoring Server on z/OS, see the *Program Directory* that comes with that product. For information on configuring the monitoring server on z/OS, see *Configuring the Tivoli Enterprise Monitoring Server on z/OS*.

Software requirements for historical data collection and the Tivoli Data Warehouse

Software requirements for storing and retrieving historical data in the Tivoli Data Warehouse are found in *IBM Tivoli Monitoring Installation and Setup Guide*.

For the warehouse proxy and Tivoli Data Warehouse, use either DB2 UDB V8.1 (and above), Microsoft SQL Server V2000, or Oracle V9.2 or V10.1. Refer to *IBM Tivoli Monitoring Installation and Setup Guide* for more information about the Tivoli Data Warehouse and the supported databases.

z/VM operating system software requirements for the OMEGAMON XE on z/VM and Linux monitoring agent

OMEGAMON XE on z/VM and Linux is supported on the following versions of the z/VM operating system:

Table 3. Supported versions of the z/VM operating system

Version of the monitoring agent	Version of z/VM operating system and the Performance Toolkit for VM			
	Version 5.2	Version 5.3	Version 5.4	Version 6.1
Version of OMEGAMON XE on z/VM and Linux				
Version 4.1.0 - Fix Pack 001	Supported	Not supported. You must install Version 4.1.0 - Fix Pack 003 or above of this monitoring agent.	Not supported. You must install Version 4.1.0 - Fix Pack 003 or above of this monitoring agent.	Not supported. You must install Version 4.1.0 - Fix Pack 003 or above of this monitoring agent.
Version 4.1.0 - Fix Pack 002	Supported	Not supported. You must install Version 4.1.0 - Fix Pack 003 or above of this monitoring agent.	Not supported. You must install Version 4.1.0 - Fix Pack 003 or above of this monitoring agent.	Not supported. You must install Version 4.1.0 - Fix Pack 003 or above of this monitoring agent.
Version 4.1.0 - Fix Pack 003	Supported	Supported - with the latest service level of the Performance Toolkit applied.	Supported - with the latest service level of the Performance Toolkit applied.	Supported - with the latest service level of the Performance Toolkit applied.
Version 4.1.1	Supported	Supported - with the latest service level of the Performance Toolkit applied.	Supported - with the latest service level of the Performance Toolkit applied.	Supported - with the latest service level of the Performance Toolkit applied.
Version 4.1.2	Supported, but with reduced functionality. See Note below.	Supported - with the latest service level of the Performance Toolkit applied.	Supported - with the latest service level of the Performance Toolkit applied.	Supported - with the latest service level of the Performance Toolkit applied.
Version 4.2.0	Supported, but with reduced functionality. See Note below.	Supported, but with reduced functionality. See Note below.	Supported - with the latest service level of the Performance Toolkit applied.	Supported - with the latest service level of the Performance Toolkit applied.

Note: To take advantage of the most recent features implemented, install the latest version of the Performance Toolkit and the latest version of the OMEGAMON XE on z/VM and Linux monitoring agent. Reduced functionality results if you run a prior version of either the Performance Toolkit or the OMEGAMON XE on z/VM and Linux monitoring agent.

Prior to z/VM v5.3, the support used by this monitoring agent was provided in the Performance Toolkit as a Web download. In z/VM v5.3, this method was replaced and this function was delivered as part of the base product. Updates to this function are now delivered as part of the normal service stream for Performance Toolkit for z/VM v5.2, z/VM v5.3 and z/VM v5.4.

For details on the relevant APARs and their corresponding PTFs for this monitoring agent, see the support page documenting the formatted output collectors of the Performance Toolkit, located at the following address:

<http://www.vm.ibm.com/related/perfkit/pksegout.html>

The following additional requirements must be met:

- The Performance Toolkit for VM is preinstalled with z/VM, V5.3 and above, with the latest available maintenance applied. Instructions for enabling, customizing, and initializing it are located in the *Program Directory* for the Performance Toolkit. It is advisable to have the latest version of the Performance Toolkit installed.
- The Performance Toolkit collects the z/VM data and writes it to a DCSS on z/VM. A DCSS must be defined prior to use by either the Performance Toolkit or prior to use by the monitoring agent. A default DCSS, called PERFOUT, is provided with the Performance Toolkit for z/VM, v5.3 and above. The size of the PERFOUT DCSS provided with z/VM should be large enough to meet the needs of your specific environment. If you have an existing DCSS, you can use the REXX executable, called FCXSEGSZ, to determine whether the DCSS you are using is large enough to suit the requirements of your environment. See “Step 2. Estimating the size of the PERFOUT DCSS” on page 42 for details on using FCXSEGSZ.

For this monitoring agent, you use the SEGOUT option of the FCONTROL MONCOLL subcommand to initiate data collection. The name of the DCSS that will be used to hold the data must be specified on this command. During agent configuration, you are prompted for the DCSS name. If you use the default name of the DCSS, type **PERFOUT**. The name used on the FCONTROL MONCOLL SEGOUT command must match the name specified during agent configuration. See the *z/VM Performance Toolkit Reference* for information on using the FCONTROL MONCOLL subcommand.

Note: The default name **PERFOUT** is used in this document to refer to the formatted output DCSS created by the FC MONCOLL SEGOUT command. If you choose to create your own with a different name, substitute your name for PERFOUT.

You can access information on the Performance Toolkit and on the latest enhancements to the Performance Toolkit as they pertain to this monitoring agent at the following Web address:<http://www.vm.ibm.com/related/perfkit/pksegout.html>

- Add the statement *MONITORRECORDS MOSTRECORDS* to the PROFILE TCPIP file on z/VM. This statement enables the collection of TCP/IP monitor data.
- The amount of data written to the PERFOUT DCSS by the SEGOUT option of the FCONTROL MONCOLL subcommand of the Performance Toolkit is determined by the CP Monitor domains that are enabled. If a domain is disabled, no data are written to the PERFOUT DCSS for that particular domain. You must enable the CP Monitor domains for which you would like to collect data. See “Enabling the CP Monitor domains” on page 40.
- Command privilege class 'E', and a connection to the *MONITOR IUCV service on z/VM, are prerequisites for using the full capabilities of the Performance Toolkit. The *MONITOR service is used by the z/VM Control program to supply monitor data to virtual machines connected to the MONITOR IUCV service. Additional information on the use of this service is available in *z/VM: CMS Application Development Guide for Assembler*.
- With the Take Action feature available with this monitoring agent, you can execute CP commands, CMS commands, and REXX execs on the z/VM operating system. To enable Take Action commands, you must create and configure a CMS guest system on the same host on which the monitoring agent resides. The CMS guest system must be set up as an unattended, disconnected service machine. The user ID should have the privileges it will require for the types of commands you want to issue by means of the Take Action commands. See “Step 11. Enabling Take Action commands (optional)” on page 53 for complete details.

Note: The Take Action feature is optional. Configuring your environment for Take Action is only required if you intend to issue CP commands or CMS commands from the Tivoli Enterprise Portal interface. Additionally, when you create a situation or edit a predefined situation, you can also issue commands using the **Action** tab of the Situation Editor. The command is sent to the system when the situation becomes true. The requirements that apply to Take Action commands apply equally to using the Action tab to issue commands that are executed on the z/VM operating system.

Supported versions of the SUSE Linux Enterprise Server operating system

The OMEGAMON XE on z/VM and Linux monitoring agent is supported on these versions of the Linux for zSeries operating system:

- SUSE Linux Enterprise Server 9 for zSeries, with Service Pack 3 or later, in 31-bit mode or in 64-bit mode
- SUSE Linux Enterprise Server 10 for zSeries, 64-bit mode
- SUSE Linux Enterprise Server 11 for zSeries, 64-bit mode

Additionally, the following requirements must be met:

- The Linux version must support the Korn shell (ksh) for installation of the monitoring agent.
- Monitored guests running under z/VM need to have the Linux Monitor Stream support enabled. This support is provided by Linux Kernel 2.4.21 or later, and 2.6 or later.
- One agent instance is needed per monitored z/VM image.
- At least one Linux guest system must be installed, configured, and running one of the supported operating systems.

You can optionally use the Take Action feature in Tivoli Enterprise Portal to issue CP commands, CMS commands, and REXX executables on the z/VM operating system. If you intend to issue Take Action commands from Tivoli Enterprise Portal, the following requirements must be met:

- `sudo`, the superuser do utility for Linux-based systems, must be available and configured on the Linux guest on which the monitoring agent is running. This allows the non-root Linux guest user to run `vmcp` without having root authority. Any working version is sufficient. `sudo` allows the monitoring agent to run with temporary root authority to send Take Action commands to the Command Processor virtual machine. See “Step 11. Enabling Take Action commands (optional)” on page 53.
- The Take Action feature requires the z/VM CP interface device driver (`vmcp`). `vmcp` is required by the monitoring agent running on Linux to issue CP commands. See “Step 11. Enabling Take Action commands (optional)” on page 53. See the *Device Drivers, Features and Commands*. See the Tivoli Enterprise Portal online help for information on using Take Action commands.

Supported versions of the Red Hat Enterprise Linux operating system

The OMEGAMON XE on z/VM and Linux monitoring agent is supported on these versions of the Red Hat Enterprise Linux operating system:

- Red Hat Enterprise Linux v.4 Update 5
- Red Hat Enterprise Linux 5

Important: For Red Hat Enterprise Linux 5, the IBM JRE for Linux on System z depends on shared libraries that are not installed by default. The installation of this monitoring agent and of the Tivoli Monitoring Agent for Linux OS (as well as Tivoli Enterprise Portal Server and Tivoli Enterprise Monitoring Server) may require that certain shared libraries be installed.

In Red Hat Enterprise Linux 5, the package managers (RPMs) that contain these libraries are:

- `libXp-1.0.0-8` (All platforms)
- `compat-libstdc++-33-3.2.3`
- `compat-libstdc++-296-2.95.3` (zSeries)

For instructions on including the shared libraries during the Red Hat Enterprise Linux 5 installation, refer to the *IBM Runtime Environment for Linux platforms, Java 2 Technology Edition, Version 1.4.2 User Guide*.

Additionally, the following requirements must be met:

- The Linux version must support the Korn shell (ksh) for installation of the monitoring agent.
- Monitored guests running under z/VM need to have the Linux Monitor Stream support enabled. This support is provided by Linux Kernel 2.4.21 or later, and 2.6 or later.

- One agent instance is needed per monitored z/VM image.
- At least one Linux guest system must be installed, configured, and running one of the supported operating systems.

You can optionally use the Take Action feature in Tivoli Enterprise Portal to issue CP commands, CMS commands, and REXX executables on the z/VM operating system. If you intend to issue Take Action commands from Tivoli Enterprise Portal, the following requirements must be met:

- `sudo`, the superuser do utility for Linux-based systems, must be available and configured on the Linux guest on which the monitoring agent is running. This allows the non-root Linux guest user to run `vmcp` without having root authority. Any working version is sufficient. `sudo` allows the monitoring agent to run with temporary root authority to send Take Action commands to the Command Processor virtual machine. See “Step 11. Enabling Take Action commands (optional)” on page 53.
- The Take Action feature requires the z/VM CP interface device driver (`vmcp`). `vmcp` is required by the monitoring agent running on Linux to issue CP commands. See “Step 11. Enabling Take Action commands (optional)” on page 53. See the *Device Drivers, Features and Commands*. See the Tivoli Enterprise Portal online help for information on using Take Action commands.

Tivoli Monitoring Agent for Linux OS monitoring agent prerequisite

OMEGAMON XE on z/VM and Linux optionally requires that the Tivoli Monitoring Agent for Linux OS agent be installed and configured. This requirement is necessary if you want to link to Tivoli Monitoring Agent for Linux OS workspaces from the OMEGAMON XE on z/VM and Linux monitoring agent. See “Dynamic linking to cross-product workspaces” on page 9. See the *IBM Tivoli Monitoring Installation and Setup Guide* for details on installing the Tivoli Monitoring Agent for Linux OS. See also the *IBM Tivoli Monitoring: Agent for Linux OS User's Guide* for information on this monitoring agent.

Note: There is no support for installing a 31-bit version of Tivoli Monitoring Agent for Linux OS on a 64-bit Linux system. Additionally, IBM Tivoli Monitoring, V6.1.0 with Fix Pack 007 and above is required for the installation and configuration of Tivoli Monitoring Agent for Linux OS.

Supported hardware

Most of the hardware required to run OMEGAMON XE on z/VM and Linux is determined by operating system considerations. For example, the requirements for the UNIX systems where parts of the IBM Tivoli Monitoring are running is determined by the operating system.

For most hardware prerequisites, consult the IBM Tivoli Monitoring documentation.

The following remaining prerequisites apply to the OMEGAMON XE on z/VM and Linux monitoring agent:

- Memory - Between 20 MB and 25 MB RAM at a minimum.
- Disk space - 175 MB for the OMEGAMON XE on z/VM and Linux monitoring agent.

If you want to link from the OMEGAMON XE on z/VM and Linux monitoring agent workspaces to the Tivoli Monitoring Agent for Linux OS workspaces, you will need to install the Tivoli Monitoring Agent for Linux OS. Together these monitoring agents take up 275 MB of disk space. See “Dynamic linking to cross-product workspaces” on page 9. For historical data disk space, see “Disk capacity planning for historical data” on page 68.

Product packaging

If you are installing the IBM Tivoli Management Services and monitoring agents for the first time, you will find familiar IBM packaging types (such as Passport Advantage[®]), installation tools (such as VMSES/E or InstallShield), and installation documentation (such as an installation guide or a program directory).

The OMEGAMON XE on z/VM and Linux monitoring agent uses both InstallShield and the VMSES/E installation tool.

The tapes provided with some of the zSeries monitoring agents are in the standard format that IBM software manufacturing uses to create the tape images for installation, such as ServerPac and Custom-Built Product Delivery Offering (CBPDO).

The following products are part of this package:

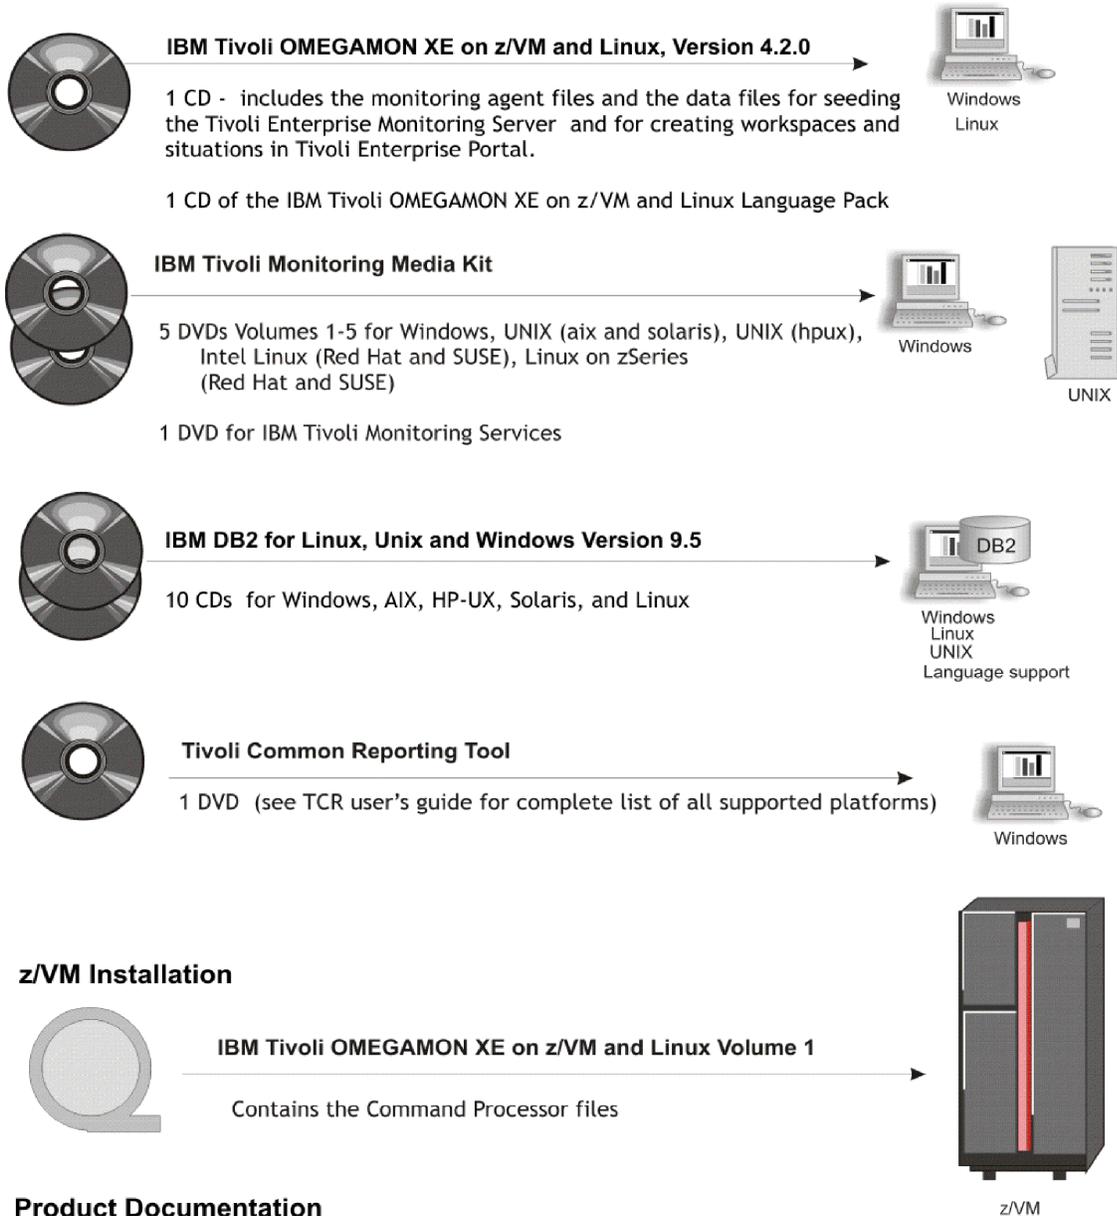
- The IBM Tivoli Management Services product and all of its components.
- The IBM OMEGAMON XE on z/VM and Linux monitoring agent.
- The IBM Tivoli Monitoring Agent for Linux OS monitoring agent is on the IBM Tivoli Management Services product CD. If you intend to use the dynamic workspace linking feature, this monitoring agent must be installed and configured. See “Dynamic linking to cross-product workspaces” on page 9.
- The Tivoli Common Reporting tool.

Each OMEGAMON zSeries product provides a program directory that describes the installation steps required to move the product code from the distribution media to your DASD, whether it is distributed on tape or electronically. See the *Program Directory* for this monitoring agent for installation instructions.

The IBM Tivoli OMEGAMON XE on z/VM and Linux product package

The contents of the IBM Tivoli OMEGAMON XE on z/VM and Linux product package are shown in Figure 4 on page 27:

Distributed Installation



Hardcopy publications

IBM Tivoli OMEGAMON XE on z/VM and Linux Program Directory
IBM Tivoli OMEGAMON XE on z/VM and Linux License Information
IBM International Program License Agreement (IPLA)
IBM Tivoli Monitoring Services on z/OS IPLA License Information
IBM Tivoli Monitoring Quick Start Guide

Figure 4. IBM Tivoli OMEGAMON XE on z/VM and Linux product packaging

Figure 4 shows the media provided with the product. This media can be divided into three groups:

Distributed Installation, which includes the following:

- *IBM Tivoli Management Services* product CDs that contain separate subdirectories and installation procedures on Windows, UNIX, Intel Linux, and Linux on zSeries operating systems for the following products and their components:
 - Tivoli Enterprise Monitoring Server Framework
 - Tivoli Enterprise Monitoring Agents (includes the required IBM Tivoli Monitoring Agent for Linux OS monitoring agent)
 - Tivoli Enterprise Portal Server Framework
 - Tivoli Enterprise Portal Desktop Client Framework
- *IBM Tivoli Management Services on z/OS Language Pack* CD
- *IBM DB2 Universal Database™* Workgroup Server Edition V9.5 CDs for Windows, AIX®, HP-UX, Solaris, and Linux, which are used by IBM Tivoli Management Services.
- *IBM Tivoli OMEGAMON XE on z/VM and Linux* CD that contains the monitoring agent files and the application support files for the monitoring agent. This CD also contains the .baroc file for this monitoring agent. The .baroc file contains the Tivoli Enterprise Console classes that are generated by IBM Tivoli Monitoring. In order to view the event data in the Tivoli Enterprise Console, you need to install this .baroc file on the event server. You add the .baroc file after you have added application support for the agent to the monitoring server.
- *IBM Tivoli OMEGAMON XE on z/VM and Linux* Language Pack CD.
- *IBM Tivoli Common Reporting Tool* DVD. The Tivoli Common Reporting tool is a reporting feature available to users of Tivoli products. This tool provides a consistent approach to viewing and administering reports. Tivoli products can provide report packages designed for use with Tivoli Common Reporting, with reports that use a consistent look and feel.

z/VM Installation that contains a tape, the *Program Directory* for performing the VMSES/E installation, and the License information. The tape contains the Command Processor files. The Command Processor enables you to send Take Action commands to z/VM from the Tivoli Enterprise Portal. This tape includes the following files:

- KVL CMD EXEC - The sample Command Processor REXX executable that serves as an example of how an executable written by you might be submitted as a Take Action from the monitoring agent. See “Step 11. Enabling Take Action commands (optional)” on page 53.
- KVL DATA contains information about the logs that were generated during the running of this monitoring agent. This log contains the following information:
 - LAST_COUNT=<number> - The last count for the number of lines that were generated.
 - LAST_LOG=<number> - The filetype for the last log that was generated.
- KVL CONFIG - A file containing a list of allowed users (AGENT_ID = the VM user ID of the Linux guest where this monitoring agent is running). It includes a list of sample commands that are not allowed, for example, CMDS=LOGOFF.

You can also specify logging the output from the processed commands to a log file. **Important:** Turning results logging on will result in significantly larger log files. Modify the log settings to meet the needs of your environment. See “z/VM requirements for Take Action commands” on page 53 for details on the settings associated with this file.

Product Documentation, which includes the following:

- Documentation CD that contains the publications for the IBM Tivoli OMEGAMON XE on z/VM and Linux monitoring agent and for the IBM Tivoli Monitoring base product.
- The Program Directory, provided in hardcopy format.
- The IBM Tivoli OMEGAMON XE on z/VM and Linux License Information, provided in hardcopy format.
- The IBM International Program License Agreement (IPLA), provided in hardcopy format.
- The IBM Tivoli Management Services on z/OS IPLA License Information, provided in hardcopy format.
- The IBM Tivoli OMEGAMON Quick Start Memo, provided in hardcopy format.

IBM Tivoli Management Services Products that are installed on Windows or UNIX systems use an installation program specific to that environment. The product tape for installing mainframe components is in CBPDO or ServerPac format.

Installation Flow

This section describes the tools you will use to install and configure OMEGAMON XE on z/VM and Linux:

- The Java-based proprietary installation tool for the monitoring agent
- InstallShield for the IBM Tivoli Management Services components
- VMSES/E for the z/VM Command Processor installation.

The installation steps required to move the monitoring agent product code from the distribution media to your environment are covered in the *Program Directory* and in the *IBM Tivoli Monitoring Installation and Setup Guide*. Additional configuration information is provided in Chapter 5, “Configuration required for the OMEGAMON XE on z/VM and Linux monitoring agent,” on page 37.

This section includes an installation flow diagram for major tasks. See Figure 5 on page 30.

About VMSES/E

The VMSES/E program is the basic tool for installing and maintaining software in z/VM systems and subsystems. This program controls changes at the element level by:

- Selecting the proper levels of elements to be installed from a large number of potential changes
- Calling system utility programs to install the changes
- Keeping records of the installed changes

VMSES/E is an integral part of the installation, service, and maintenance processes for z/VM software products and product packages. The guidance for performing a VMSES/E installation is the *Program Directory* for this product.

Installation flow

The basic flow of high-level installation and configuration tasks is shown in Figure 5 on page 30.

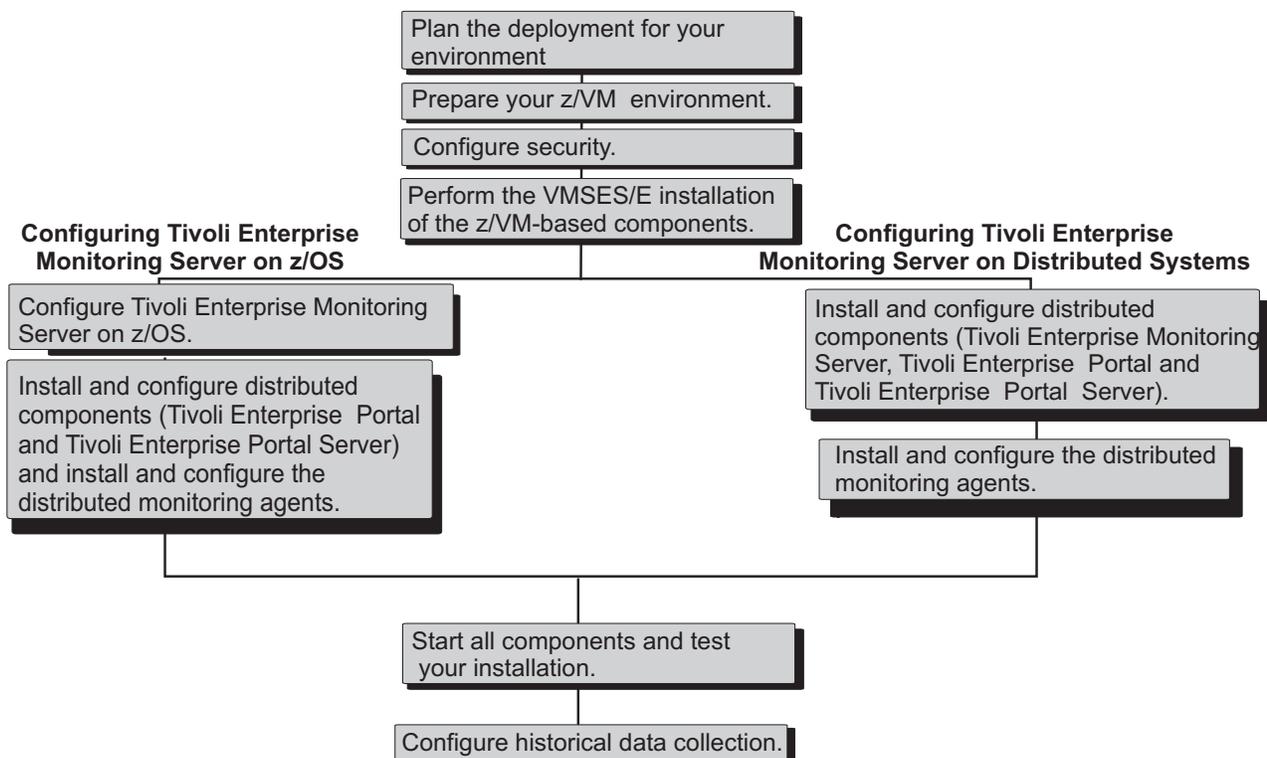


Figure 5. Installation and configuration overview

Planning the deployment for your environment

The IBM Tivoli Monitoring environment requires installation and configuration to be performed on the managed systems as well as on the distributed systems where some of the components run.

Before you proceed with installation:

- Choose which systems to monitor.
- Decide where you will deploy the hub Tivoli Enterprise Monitoring Server, the Tivoli Enterprise Portal Server, and Tivoli Enterprise Portal.
- Decide if and where you will deploy remote Tivoli Enterprise Monitoring Servers.
- Complete the planning worksheets in “Planning worksheets” on page 79.

Planning information is also available for the IBM Tivoli Monitoring components, in the locations shown below:

- Planning for the distributed components of the IBM Tivoli Monitoring is found in the *IBM Tivoli Monitoring Installation and Setup Guide*.
- Planning for a Tivoli Enterprise Monitoring Server on z/OS is found in *Configuring IBM Tivoli Enterprise Monitoring Server on z/OS*.
- Planning for the Tivoli Data Warehouse and historical data collection is found in the *IBM Tivoli Monitoring Administrator's Guide* and in the *IBM Tivoli Monitoring Installation and Setup Guide*.

Preparing your z/VM environment

To make the OMEGAMON XE on z/VM and Linux environment fully operational, you must perform some z/VM environment set up tasks. These tasks include enabling the Performance Toolkit and making sure it is collecting data. The tasks also include installing and configuring the Command Processor virtual machine and ensuring that it is activated.

This step does not have to be performed before you perform your monitoring agent installation and configure the components, but it must be completed before you bring IBM Tivoli Monitoring up and start the components.

The information required to proceed with the z/VM preparation is found in Chapter 5, “Configuration required for the OMEGAMON XE on z/VM and Linux monitoring agent,” on page 37.

Configuring security

Before you enable security in the IBM Tivoli Monitoring, you must define the user ID and password for each user who will logon to the Tivoli Enterprise Portal. The security definitions will be performed on the system where you run the hub Tivoli Enterprise Monitoring Server.

This step does not have to be performed before you perform your installation and configure the components, but it must be completed before you bring IBM Tivoli Monitoring up and start the components.

The information required to do this security configuration is found in the *IBM Tivoli Monitoring Installation and Setup Guide*.

Configuring the Tivoli Enterprise Monitoring Server on a distributed operating system

If your Tivoli Enterprise Monitoring Server is on a distributed operating system, the configuration sequence is as follows:

1. Install and configure the distributed components. The following components must be installed on the systems identified in your planning worksheet:

- Tivoli Enterprise Monitoring Server
- Tivoli Enterprise Portal Server
- Tivoli Enterprise Portal desktop client

Note: When configuring the monitoring server, be sure to perform the step that adds application support to the monitoring server. The monitoring server must be started to perform this task.

2. Install and configure the monitoring agents.

- Install and configure the OMEGAMON XE on z/VM and Linux monitoring agent.
- Install and configure the Tivoli Monitoring Agent for Linux OS monitoring agent. This is only required if you intend to use the “Dynamic linking to cross-product workspaces” on page 9 feature.

For information on installing and configuring the OMEGAMON XE on z/VM and Linux monitoring agent, refer to the *IBM Tivoli Monitoring Installation and Setup Guide*. For additional configuration information for this monitoring agent, see Chapter 5, “Configuration required for the OMEGAMON XE on z/VM and Linux monitoring agent,” on page 37.

For information on installing the IBM Tivoli Monitoring Agent for Linux OS monitoring agent, refer to the *IBM Tivoli Monitoring Installation and Setup Guide* and to the *IBM Tivoli Monitoring: Agent for Linux OS User's Guide*.

Configuring the Tivoli Enterprise Monitoring Server on z/OS

If you have chosen to run your hub Tivoli Enterprise Monitoring Server on z/OS, it is assumed that you have already installed the Tivoli Enterprise Monitoring Server on z/OS using the Configuration Tool. See *Configuring Tivoli Enterprise Monitoring Server on z/OS*.

Starting and verifying the components

After you have installed and configured all components required to monitor network performance with OMEGAMON XE on z/VM and Linux, start the components in the required order and verify the configuration.

Refer to Chapter 5, “Configuration required for the OMEGAMON XE on z/VM and Linux monitoring agent,” on page 37 for information about performing this step. See also the *IBM Tivoli Monitoring Installation and Setup Guide*.

Preparing to create historical reports

The tasks required to enable historical reporting can be performed at any time after you have verified your installation. Historical data are organized into short-term (by default, up to the last 24 hours) and long-term. You can configure short-term historical data to be on the system where the monitoring agent resides or on the system where the Tivoli Enterprise Monitoring Server resides.

For detailed information about collecting historical data, see the *IBM Tivoli Monitoring Administrator's Guide*.

Chapter 4. Upgrading IBM Tivoli Monitoring

This chapter contains some topics to consider if you currently have one or more monitoring agents in your environment.

You will need to upgrade the Tivoli Enterprise Portal, Tivoli Enterprise Monitoring Server, Tivoli Data Warehouse, and the IBM Tivoli Monitoring Agent for Linux OS monitoring agent.

Note: You only need to upgrade the IBM Tivoli Monitoring Agent for Linux OS if you intend to use the dynamic workspace linking feature. See “Dynamic linking to cross-product workspaces” on page 9.

Upgrading an existing IBM Tivoli Monitoring environment

If you have an existing installation of IBM Tivoli Monitoring, you may need to upgrade to newer versions of Tivoli Enterprise Portal and IBM Tivoli Monitoring components. Identify the systems where you will need to upgrade or install Tivoli Enterprise Portal and IBM Tivoli Monitoring components.

The following lists the components and the minimum required versions of IBM Tivoli Monitoring for this monitoring agent:

- Tivoli Enterprise Portal V6.2.1 with Interim Feature 004 and above
- Tivoli Enterprise Portal Server V6.2.1 with Interim Feature 004 and above
- Tivoli Enterprise Monitoring Server V6.2.1 with Interim Feature 004 and above

See the *IBM Tivoli Monitoring Installation and Setup Guide* for installation instructions.

Installation in a new environment

If you do not have an existing installation you must install version 6.2.1 with Interim Feature 004 or later of IBM Tivoli Monitoring and the required components listed above. See the *IBM Tivoli Monitoring Installation and Setup Guide* for installation instructions.

Part 2. Configuration required for this monitoring agent

The chapters in this section walk you through the process of configuration required for the OMEGAMON XE on z/VM and Linux monitoring agent.

- Chapter 5, “Configuration required for the OMEGAMON XE on z/VM and Linux monitoring agent,” on page 37 lists the requirements that must be met before configuring this agent. This chapter also provides the configuration steps that must be performed.
- Chapter 6, “Defining user IDs and security,” on page 61 provides information about user IDs and how user security is implemented in the monitoring agents.

Chapter 5. Configuration required for the OMEGAMON XE on z/VM and Linux monitoring agent

This chapter makes the following assumptions:

- You have installed IBM Tivoli Management Services and all of its components, including:
 - A Tivoli Enterprise Monitoring Server
 - A Tivoli Enterprise Portal Server
 - The Tivoli Enterprise Portal clients
 - The OMEGAMON XE on z/VM and Linux monitoring agent. Be sure to install and enable the application support for this monitoring agent. Application support adds required support files, such as situations, to the monitoring server. The application support files for this monitoring agent are on the *IBM Tivoli OMEGAMON XE on z/VM and Linux* product CD.
 - The IBM Tivoli Monitoring Agent for Linux OS monitoring agent. Be sure to install and enable the application support for the monitoring agent. Application support adds required support files, such as situations, to the monitoring server. Installation of this monitoring agent is required only if you intend to use the “Dynamic linking to cross-product workspaces” on page 9 feature.
 - The Eclipse Help Server (with this feature, you can search the provided help files for specific text strings). You can also view the help system for monitoring agents and for IBM Tivoli Monitoring. The Eclipse Help Server is installed with IBM Tivoli Monitoring.

See the *IBM Tivoli Monitoring Installation and Setup Guide*, for instructions on installing the IBM Tivoli Management Services and all of its components. That guide also provides instructions for installing this monitoring agent, and it provides instructions for installing and enabling application support for the monitoring agent on the shared technology components.

- You have installed either the z/VM operating system, v6.1 with the latest service level of the Performance Toolkit applied, or you have installed z/VM operating system, v5.4 with the latest service level of the Performance Toolkit applied. See “z/VM operating system software requirements for the OMEGAMON XE on z/VM and Linux monitoring agent” on page 22 for all of the supported versions of z/VM. See also the general z/VM documentation for details at <http://www.vm.ibm.com/library/>.
- You have configured and activated the Performance Toolkit for VM, with the latest maintenance installed. The Performance Toolkit for VM is pre-installed on z/VM. Instructions for enabling, customizing, and initializing it are located in the *IBM Tivoli OMEGAMON XE on z/VM and Linux Program Directory* for the Performance Toolkit. See also the *z/VM: Performance Toolkit Guide* and *z/VM: Performance Toolkit Reference*.

You can access information on the Performance Toolkit and on the latest enhancements to the Performance Toolkit as they pertain to this monitoring agent at the following Web address:

<http://www.vm.ibm.com/related/perfkit/pksegout.html>

- During configuration of this monitoring agent, you are prompted to enter the name of the DCSS. If you use the default PERFOUT DCSS, press **Enter**. Otherwise, enter the name of the DCSS that you defined. For more information, see “Configuration steps” on page 40.

During configuration, you are also prompted to enter the name of the z/VM user ID to be used to process Take Action commands. This is optional and required only if you intend to issue Take Action commands on the z/VM system. Press **Enter** if you do not intend to issue Take Action commands. See “Step 11. Enabling Take Action commands (optional)” on page 53 for additional information on enabling Take Action commands.

If you’re installing as a non-root user, you will be prompted for the root password to update the auto start script. If you have the root password, enter it now. If you don’t have the root password, press **Enter** when prompted for the root password. You will get an error message stating that the system is unable to update the auto restart script. Updating the auto restart script requires root authority. If you want the agent to start up automatically when the system is initialized, the root password is required.

- You have enabled the CP Monitor service so that the Performance Toolkit can receive the data. See the z/VM documentation for information on enabling the CP Monitor service.
- The amount of data written to the PERFOUT DCSS by the SEGOUT option of the FCONTROL MONCOLL subcommand of the Performance Toolkit is determined by the CP Monitor domains that are enabled. If a domain is disabled, no data are written to the PERFOUT DCSS for that type of data. You must enable the CP Monitor domains for which you would like data to be collected. See “Enabling the CP Monitor domains” on page 40.
- You have installed the Command Processor. This step will have been performed if you want to use the Take Action feature. See the *Program Directory* for instructions. See also “Step 11. Enabling Take Action commands (optional)” on page 53.
- If you intend to use the dynamic workspace linking feature, the IBM Tivoli Monitoring Agent for Linux OS monitoring agent must be installed, configured, and active. See the *IBM Tivoli Monitoring: Agent for Linux OS User's Guide* for details on this monitoring agent. See the *IBM Tivoli Monitoring Installation and Setup Guide* for installation instructions. See also “Dynamic linking to cross-product workspaces” on page 9.

Important: Once you have installed and configured all of the necessary components, see “Required order of tasks for viewing data at the monitoring agent” for a list of the high-level tasks that are required to view data, in the order in which they must be performed.

See Chapter 3, “Installation preparation,” on page 21 for the list of software and hardware prerequisites.

To obtain the most recent installation updates for the OMEGAMON XE on z/VM and Linux monitoring agent, review the *Readme* information for this product.

Required order of tasks for viewing data at the monitoring agent

This section assumes that all of the required components have been installed and configured. Review the “Planning worksheets” on page 79 and the configuration steps in this chapter. The table that follows lists the high-level tasks that are required to view data at the monitoring agent, in the order in which they must be performed.

Table 4. Required order of tasks for viewing data at the monitoring agent

Task	Description	Who performs the task	Where to find information
1	Start the Tivoli Enterprise Monitoring Server.	IBM Tivoli Monitoring administrator	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>
2	Start DB2 on the Tivoli Enterprise Portal Server, if not already started.	IBM Tivoli Monitoring administrator	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>
3	Start the Tivoli Enterprise Portal Server.	IBM Tivoli Monitoring administrator	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>
4	Start the Tivoli Data Warehouse (optional - only if you intend to collect historical data).	IBM Tivoli Monitoring administrator	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>
5	Verify that the application support for this monitoring agent has been installed and enabled on the appropriate IBM Tivoli Monitoring shared components. The application support files for this agent are on the product CD.	IBM Tivoli Monitoring administrator	<i>IBM Tivoli Monitoring Installation and Setup Guide</i> and <i>Configuring IBM Tivoli Enterprise Monitoring Server on z/OS</i>
6	Enable the CP Monitor domains for the type of data that you want to collect.	z/VM system administrator	“Configuring TCP/IP on z/VM” on page 40 and “Enabling the CP Monitor domains” on page 40

Table 4. Required order of tasks for viewing data at the monitoring agent (continued)

Task	Description	Who performs the task	Where to find information
7	Issue the QUERY NSS NAME PERFOUT command to ensure that the PERFOUT DCSS is defined and available (if you are using the default name of the DCSS).	z/VM system administrator	<i>z/VM: CP Commands and Utilities Reference</i>
8	Start the Performance Toolkit and issue the FC MONCOLL SEGOUT ON PERFOUT command.	z/VM system administrator	<i>z/VM: Performance Toolkit Reference</i> and “Step 4. Issue the FC MONCOLL SEGOUT ON PERFOUT command” on page 47
9	Start the optional Command Processor virtual machine (if not already started). The Command Processor is used for Take Action commands.	z/VM system administrator	Refer to the following: <ul style="list-style-type: none"> • <i>IBM Tivoli OMEGAMON XE on z/VM and Linux Program Directory</i> for instructions on installing this executable file. • The z/VM Installation section of “The IBM Tivoli OMEGAMON XE on z/VM and Linux product package” on page 26 for the contents of the tape that contains the Command Processor file. • “Step 11. Enabling Take Action commands (optional)” on page 53.
10	Load the DCSS driver on the Linux on zSeries guest system where the z/VM agent will be running.	Linux system administrator	“Step 6. Loading the DCSS device driver” on page 48 and “Step 7. Adding the PERFOUT DCSS to your Linux guest” on page 49
11	Start the application monitor data collection for each Linux guest.	Linux system administrator	“Step 9. Enabling the collection of Linux data” on page 49
12	Start the OMEGAMON XE on z/VM and Linux monitoring agent on the Linux on zSeries guest system to be monitored.	Linux system administrator	“Starting the monitoring agent” on page 58
13	Start the Tivoli Monitoring Agent for Linux OS on the Linux guests that you want to monitor. This is optional, but required if you intend to use the dynamic workspace linking feature.	Linux system administrator	“Starting the monitoring agent” on page 58 and “Step 11. Enabling Take Action commands (optional)” on page 53
14	Log on to the Tivoli Enterprise Portal and observe the collected data. It may be necessary to allow two sample intervals to elapse. Use the z/VM PTK Collector Status view of the z/VM Linux Systems workspace to verify that the Performance Toolkit is active, and that data is being collected for the types of data that are enabled.	IBM Tivoli Monitoring administrator	<i>IBM Tivoli Monitoring Installation and Setup Guide</i> and the <i>IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide</i> or the online help system for the monitoring agent

Configuration steps

After you have installed all of the components as outlined in the introduction to this chapter, you are ready to perform the configuration steps required for this monitoring agent. Perform the following steps in this order:

- “Step 1. Enabling the collection of data”
- “Step 2. Estimating the size of the PERFOUT DCSS” on page 42
- “Step 3. Defining a DCSS on z/VM” on page 45
- “Step 4. Issue the FC MONCOLL SEGOUT ON PERFOUT command” on page 47
- “Step 5. Configuring the DCSS device driver on the Linux guest” on page 47
- “Step 6. Loading the DCSS device driver” on page 48
- “Step 7. Adding the PERFOUT DCSS to your Linux guest” on page 49
- “Step 8. Loading the PERFOUT DCSS at startup time” on page 49
- “Step 9. Enabling the collection of Linux data” on page 49
- “Step 10. Enabling dynamic workspace linking” on page 52
- “Step 11. Enabling Take Action commands (optional)” on page 53
- “Step 12. Installing the Language Packs (optional)” on page 57

Step 1. Enabling the collection of data

Several configuration steps are required to enable the collection of data for this monitoring agent. These configuration steps are outlined in the sections that follow.

Configuring TCP/IP on z/VM

Perform the following steps to configure TCP/IP on z/VM:

1. The MONITORRECORDS MOSTRECORDS statement enables the collection of TCP/IP monitor data. Add the following statement to the PROFILE TCPIP file on z/VM:

```
MONITORRECORDS MOSTRECORDS
```

If your PROFILE TCPIP file contains the MONITORRECORDS ALL RECORDS statement, you do not need to update the PROFILE TCPIP file. However, be aware that the MONITORRECORDS ALLRECORDS statement yields monitor data that is not used by the monitoring agent, and could result in excessive data collection.

2. For the MONITORRECORDS ALL RECORDS statement to take effect, the TCPIP virtual machine must be authorized to create monitor data records. Add the OPTION APPLMON statement to the User Directory entry for the TCPIP virtual machine.
3. To collect monitor data, enable the APPLDATA class for both SAMPLE and EVENT recording for the TCPIP virtual machine.
4. To ensure that TCP/IP configuration information is included in the monitor data file, start the monitor and a monitor writer before initializing the TCPIP virtual machine.

For details on configuring TCP/IP, see the *TCP/IP Level 3A0 Planning and Customization* guide.

Enabling the CP Monitor domains

The data records written to the PERFOUT DCSS are based on the types of records with which they are associated. You must enable the CP Monitor domains for each type of data to be collected. This step must be performed on the z/VM system where the Performance Toolkit is running. If a domain is disabled, no data are written to the PERFOUT DCSS for that particular domain. The use of the CP MONITOR SAMPLE and EVENT commands requires privilege class 'E'.

The table that follows lists the types of data that can be written to the PERFOUT DCSS and displayed in the Tivoli Enterprise Portal for this monitoring agent. The MONITOR SAMPLE and MONITOR EVENT domains that must be enabled for each type of data are included in this table.

Table 5. CP Monitor domains that must be enabled

Data record type	Associated CP Monitor domain
Channels	MONITOR SAMPLE domain: SYSTEM
CCW Translations	MONITOR SAMPLE domain: SYSTEM
CP-owned minidisks	MONITOR SAMPLE domain: STORAGE
Cached control units	MONITOR SAMPLE domain: I/O
DASD cache	MONITOR SAMPLE domain: I/O
Defined LPARs	MONITOR SAMPLE domain: SYSTEM
FICON® channels	MONITOR SAMPLE domain: SYSTEM
HiperSockets™ devices	MONITOR SAMPLE domain: SYSTEM
Linux guest systems	MONITOR SAMPLE domain: APPLDATA and MONITOR EVENT domain: APPLDATA. See also "Step 9. Enabling the collection of Linux data" on page 49.
Logged-on users	MONITOR SAMPLE domain: USER
Minidisk cache	MONITOR SAMPLE domain: SYSTEM
Monitored disk devices	MONITOR SAMPLE domain: I/O
Network virtual switch devices	MONITOR SAMPLE domain: I/O
Network server virtual machines	MONITOR SAMPLE domain: APPLDATA and MONITOR EVENT domain: APPLDATA
Network users	MONITOR SAMPLE domain: APPLDATA and MONITOR EVENT domain: APPLDATA
Paging and spooling areas	MONITOR SAMPLE domain: STORAGE
Processor	MONITOR SAMPLE domain: SYSTEM and MONITOR SAMPLE domain: PROCESSOR
Real storage	MONITOR SAMPLE domain: SYSTEM
Resource constraint	MONITOR SAMPLE High Frequency Sampling Active
Spin Lock	MONITOR SAMPLE domain: SYSTEM
System	MONITOR SAMPLE domain: SYSTEM
System Health	MONITOR SAMPLE domain: SYSTEM
TCP/IP	MONITOR SAMPLE domain: APPLDATA
TCP/IP users	MONITOR EVENT domain: APPLDATA
Virtual disk	MONITOR SAMPLE domain: STORAGE
Workload	MONITOR EVENT domain: APPLDATA

High-frequency sample data is a set of counters and data that represent the state of the system at the moment they are sampled. These samples are the source of data reported on the Resource Constraint workspace. High-frequency sampling is automatically activated whenever the MONITOR SAMPLE START command is issued, unless the RATE value has been previously set to STOP. Active high-frequency sampling in the CP Monitor is a prerequisite for Resource Constraint workspace views.

The following examples show how you would enable the CP Monitor SAMPLE APPLDATA domain and the CP Monitor EVENT APPLDATA domain:

- To enable the CP Monitor to collect SAMPLE APPLDATA, enter this command on the z/VM system where the Performance Toolkit is running:

```
CP MONITOR SAMPLE ENABLE APPLDATA ALL
```
- To enable the CP Monitor to collect EVENT APPLDATA, enter this command on the z/VM system where the Performance Toolkit is running:

```
CP MONITOR EVENT ENABLE APPLDATA ALL
```

The Performance Toolkit collects all the data it can use. This can cause considerable load if you have enabled EVENT data for all users on a large system: records will be generated by CP, and collected by Performance Toolkit, for each single user transaction. Care should be taken. See the *z/VM Performance Toolkit* guides.

See the *z/VM: CP Commands and Utilities Reference* for details on the MONITOR SAMPLE and on the MONITOR EVENT commands and the operands associated with each command.

Step 2. Estimating the size of the PERFOUT DCSS

A DCSS is a shared memory area outside of (or discontinuous from) the address range of a VM user ID. CP allows multiple VM user IDs to access the shared memory area. Each VM user ID accesses the shared memory area at the same address range. Only one copy of the contents of a DCSS is required and can be shared by all VM user IDs.

The product-provided DCSS, called **PERFOUT**, should be large enough to accommodate the requirements of your environment. If you create your own DCSS, you must determine the approximate size of the DCSS, based upon how many devices are in use, how many users are accessing the system, and how many items are being monitored. For a small to medium size system, the DCSS must be able to contain at least 6 MB of storage. In number of pages, this translates to 600 hexadecimal pages, or 1536 decimal pages.

You can estimate the size of the DCSS that you'll need by using the FCXSEGSZ REXX executable provided by the Performance Toolkit. This executable runs on any virtual machine where the CMS is running. FCXSEGSZ does not delete any segments and it does not define any segments. It is simply a size calculator. You need to provide input each time you invoke the tool, as the input is not saved across invocations.

FCXSEGSZ calculates the size of the DCSS based upon the information that you provide. Your responses to the prompts when running this executable must take into account normal fluctuations in system utilization in your installation. Your responses must also take into account expected increases in defined resources or increases in numbers of users. It is best to provide higher counts. Providing higher counts for items in your enterprise saves you from having to increase counts in the future. The contents of a DCSS are not saved in the spool file; only the definitions for the DCSS are saved, thus the DCSS itself occupies a minor amount of spool space.

The size of the DCSS is based on numerous distinct record types (listed below). The data records written to the DCSS vary in length based on the type of record with which they are associated. The frequency of data collection is tied to the CP Monitor SAMPLE interval setting. The monitor SAMPLE interval is used for all the permanently collected data (the Performance Toolkit uses this interval also when obtaining data for any of the general performance displays directly from CP control blocks).

Note: The amount of data written to the DCSS by the SEGOUT option of the FCONTROL MONCOLL subcommand of the Performance Toolkit is determined by the CP Monitor domains that are enabled. If a domain is disabled, no data are written to the DCSS for that particular domain. You must enable the CP Monitor domains for the types of data to be collected. See “Enabling the CP Monitor domains” on page 40. See the *z/VM: Performance Toolkit Reference* for information on using the FCONTROL MONCOLL subcommand.

To run FCXSEGSZ, type the following command at the CMS virtual machine that has access to the disk where the Performance Toolkit is installed:

```
FCXSEGSZ
```

When you run FCXSEGSZ, you are prompted to provide counts for the following items in your enterprise:

1. Defined LPARs
2. Monitored disk devices
3. Paging and SPOOLing areas
4. Logged-on users
5. HiperSockets devices

6. Network virtual switch devices
7. Network server virtual machines
8. Network users
9. Linux virtual machines
10. LPAR processor engines
11. z/VM channels
12. LPAR channels
13. FICON channels
14. Cache DASD devices
15. Cache control units
16. Virtual disks in storage (VDISKS)

To skip a count because you do not monitor a type of item, specify **1** at the prompt for that item.

1. Defined LPARs

Enter the number of LPAR partitions that are currently defined or that will be defined in the processor complex where the monitored z/VM system resides. Provide a count of all partitions defined, regardless of the operating system. Include the monitored z/VM partition in the count. Type a value from 1 to 255.

2. Monitored disk devices

Enter the number of disk devices that are attached to the monitored z/VM system. If you use the MONITOR SAMPLE commands to limit the number of I/O devices monitored, enter the maximum number of devices that are enabled to MONITOR SAMPLE. Type a value from 1 to 65535.

Note: Do not include SCSI devices in this count.

3. Paging and SPOOLing areas

Enter the maximum number of paging and spooling areas that are expected to be online on the monitored z/VM system. An area is defined as a single contiguous allocation of cylinders (CKD devices) or blocks (FBA devices) defined to CP as paging or spooling space. If a device has multiple paging or spooling areas, each separate area must be counted. Type a value from 1 to 255.

4. Logged on users

Enter the maximum number of users logged on concurrently to the monitored z/VM system. Type a value from 1 to 999999.

5. HiperSockets devices

Enter the maximum number of HiperSockets channel paths online on the monitored z/VM system. If you use the MONITOR SAMPLE commands to limit the number of I/O devices monitored, enter the maximum number of HiperSockets devices that are enabled to MONITOR SAMPLE. Type a value from 1 to 65535.

6. Network virtual switch devices

Enter the maximum number of virtual switch devices defined to the monitored z/VM system. If you use the MONITOR SAMPLE commands to limit the number of I/O devices monitored, enter the maximum number of virtual switch devices that are enabled to MONITOR SAMPLE. Type a value from 1 to 65535.

7. Network server virtual machines (TCP/IP)

Enter the maximum number of network server virtual machines logged on to the monitored z/VM system. Each network server virtual machine must be configured to handle network traffic (for example, TCP/IP), and must have application monitor data collection enabled. Type a value from 1 to 255.

8. Network users

Enter the maximum number of users of the network server virtual machines specified in Step 7, that are logged on to the monitored z/VM system. The count of logged on network users includes the number of users logged on to z/VM using CMS. The count also includes the number of users logged on to network applications, such as FTP. Type a value from 1 to 999999.

9. **Linux virtual machines**

Enter the maximum number of virtual machines running as Linux on zSeries guest systems on the monitored z/VM system. Type a value from 1 to 999999.

Note: The collection of Linux data requires that you enable the MONITOR SAMPLE APPLDATA domain. See “Step 9. Enabling the collection of Linux data” on page 49. Also, the Linux Monitor Stream Support must be installed and activated within Linux.

10. **Engines defined to LPARs**

Enter the total number of processors of all types (CP, IFL, ZIIP, ZAAP, and ICF) defined to the logical partitions in the processor complex. Include all processors, either shared or dedicated, regardless of the operating system running in the owning partition. Be sure to include the processors used by the monitored z/VM logical partition in the count. Type a value from 1 to 65535.

11. **Channels defined to z/VM**

Enter the total number of channels defined to the z/VM system being monitored. Do not include FICON channels in the count. FICON channels are counted separately. Type a value from 1 to 1024.

12. **LPAR channels defined to z/VM**

Enter the total number of channels defined to the logical partition hosting the z/VM system being monitored. Do not include FICON channels in the count. FICON channels are counted separately. Type a value from 1 to 1024.

13. **FICON channels defined to z/VM**

Enter the total number of FICON channels defined to the logical partition hosting the z/VM system being monitored. Type a value from 1 to 1024.

14. **Disk cache devices defined to z/VM**

Enter the total number of disk devices with cache capability that are attached to the monitored z/VM system. If you use the CP MONITOR SAMPLE commands to limit the number of I/O devices being monitored, enter the maximum number of disk cache devices that are enabled to MONITOR SAMPLE. Type a value from 1 to 65535.

15. **Cache control units defined to z/VM**

Enter the total number of control units with cache capability that are attached to the monitored z/VM system. If you use the CP MONITOR SAMPLE commands to limit the number of I/O devices monitored, enter the maximum number of cache control units that are enabled to MONITOR SAMPLE. Type a value from 1 to 65535.

16. **VDISKS defined to z/VM**

Enter the total number of virtual disks in storage (VDISKS) that are attached to users logged on to the monitored z/VM system. Type a value from 1 to 999999.

After entering values for each of the above listed items, FCXSEGSZ runs. The following lines are sample output:

```
Minimum SEGOUT segment allocation needed (bytes, in decimal): 776848
Minimum SEGOUT segment allocation needed (pages, in decimal): 191
Minimum SEGOUT segment allocation needed (bytes, in hexadecimal): BDA90
Minimum SEGOUT segment allocation needed (pages, in hexadecimal): BF
```

The sample output displays the minimum size of the DCSS. From here you can proceed to the next step of determining the start and end address of the DCSS.

Step 3. Defining a DCSS on z/VM

The product-provided PERFOUT DCSS is large enough to fit the requirements of most environments. PERFOUT is provided with a predefined location.

If you use PERFOUT, or if you already have a DCSS defined for your environment, you may skip the first part of this step. However, you will still need to save the PERFOUT DCSS prior to use (see Step 2 below).

Note: The DCSS must not overlap the CP Monitor DCSS (default name is MONDCSS) or any of the segments used by CMS.

The DCSS is defined in the following basic format:

```
DEFSEG PERFOUT <hexpage1 - hexpage2> <type>
```

where:

- <hexpage1 - hexpage2> = the range of pages that are to be saved. Here you specify the starting page and the ending page.
- <type> = the type of virtual machine access permitted to pages in the range. For this monitoring agent, use either SN (shared read/write access, no data saved) or SW (shared read/write access). You may prefer to use SN since it incurs less overhead.

Additional options are described in the *z/VM: CP Commands and Utilities Reference*.

Use the sample output provided by running FCXSEGSZ to define the size of the DCSS. The segment allocation estimate of pages in hexadecimal is very helpful, as in the following example:

```
Minimum SEGOUT segment allocation needed (pages, in hexadecimal): 34
```

If, for example, your environment requires a segment allocation of 34 hexadecimal pages, and your starting location is x'10000000, you define the DCSS as follows:

1. Type this command on the z/VM system where the Performance Toolkit is installed:

```
DEFSEG PERFOUT 10000-10033 SN
```

where:

- PERFOUT = the default name of the DCSS for the SEGOUT option of the FCONTROL MONCOLL subcommand of the Performance Toolkit
- 10000-10033 = the range of available hexadecimal pages that are to be saved
- SN = the type of virtual machine access permitted to pages in the range. In this case, shared read/write access, with no data saved, is specified.

Note: Before you save the DCSS, make sure that the virtual machine is large enough to contain the segment. That is, the DCSS must reside entirely within the virtual machine. Also, the CMS and the MONDCSS segments must not overlap with the PERFOUT DCSS.

2. Save the PERFOUT DCSS, as follows:

```
SAVESEG PERFOUT
```

For more information on defining a DCSS, see the *z/VM: CP Commands and Utilities Reference* guide. See also the *Saved Segments Planning and Administration* guide.

Helpful tips when defining your DCSS on z/VM

The following tips can be helpful when defining your own DCSS on z/VM:

Tips:

- Use the QUERY NSS ALL command to display a list of all of the segments already defined to your z/VM system.
- Use the QUERY NSS ALL MAP to display the size and address ranges of existing segments. Make note of the beginning address and the ending address of a segment so that the addresses of the DCSS that you define do not overlap with those of existing segments.
- Use the QUERY VIRTUAL STORAGE command to display the size of storage accessible to your virtual machine.
- Use the DEFSEG command to define the DCSS. When you use this command, you also specify the name of the DCSS that will be used to hold the data. The default name of the DCSS for the SEGOUT option is PERFOUT. See the *z/VM: Performance Toolkit Guide* for details on the FCONTROL MONCOLL subcommand.
- Use the DEFINE STORAGE CONFIG command to define the storage size of the guest virtual machine.
- Use the SAVESEG command to save the defined DCSS. You cannot access a DCSS unless it has been saved.

Note: Changing the virtual machine storage size requires a re-IPL of the virtual machine.

The DEFSEG, the SAVESEG, and the QUERY NSS commands require privilege class 'E'. The QUERY VIRTUAL STORAGE and the DEFINE STORAGE commands require privilege class 'G'. For detailed information on these CP commands and on using DEFSEG to define a DCSS, see the *z/VM: CP Commands and Utilities Reference* guide. See also the *Saved Segments Planning and Administration* guide.

Running out of memory in the DCSS

The FCXSEGSZ executable calculates enough segment space to hold two data samples from the Performance Toolkit. This allows the monitoring agent sufficient time to process the most recent sample before the space is reused. During Performance Toolkit initialization, the segment is divided roughly in half, with each half holding one entire sample.

The areas used to hold sample data are allocated only when a sample is collected. There are no fixed allocations by record type. For example, during one sample, the USER data may need 14 KB of the segment, and DEVICE data will need 24 KB. In the next sample, changes to device configuration and additional users may result in 20 KB for USER data and 18 KB for DEVICE data. In either case, 38 KB of the segment will be consumed for USER data and for DEVICE data combined.

This design provides the most effective use of the allocated segment. It is still possible to allocate a segment that is not large enough to contain all the data available and enabled for collection. If the Performance Toolkit detects that a segment is not large enough, it issues the following message at the Performance Toolkit console:

```
FCXxxx774E Insufficient space in SEGOUT segment. Larger segment needed
```

This error causes SEGOUT data collection to stop. Use FCXSEGSZ to re-estimate the size of the segment. The amount of space in the segment needed by the data collectors depends on the number of users, the number of devices, and so on being monitored in your z/VM environment. After you expand the size of the PERFOUT DCSS, you must then manually restart SEGOUT data collection. See the *z/VM: Performance Toolkit Reference* for additional details on this message.

Note: When expanding the size of the PERFOUT DCSS, verify that the Linux definitions align with the segment definitions.

Step 4. Issue the FC MONCOLL SEGOUT ON PERFOUT command

At the Performance Toolkit, issue this command:

```
FC MONCOLL SEGOUT ON PERFOUT
```

The above command assumes that you are using PERFOUT as the default name of your DCSS. Otherwise, substitute the name of your DCSS. Add this command to your FCONX \$PROFILE file.

Step 5. Configuring the DCSS device driver on the Linux guest

Before you can configure the DCSS device driver on the Linux guest, you must either be using the default PERFOUT DCSS or you must have defined your own DCSS on z/VM and know the name assigned to that DCSS on z/VM.

If you use a watchdog device driver, turn off the watchdog before adding or saving a DCSS. Adding or saving a DCSS can result in a watchdog timeout, if it is active.

Determining the start and end addresses of the PERFOUT DCSS

Before you can configure the PERFOUT DCSS on the Linux guest, you first need to know the start and end addresses of the PERFOUT DCSS defined on z/VM. You can find out this information by issuing the following CP command from a CMS session with privilege class 'E':

```
QUERY NSS MAP NAME PERFOUT
```

The output gives you the start and end addresses of the defined PERFOUT DCSS in units of 4-KB pages. Your guest storage must not overlap with the address range of the PERFOUT DCSS, and it must match up with the location of the DCSS.

Depending upon the start and end addresses of the PERFOUT DCSS and the size of the Linux guest memory, you can use either one of these two methods:

- Define the guest storage as two or more discontinuous storage extents such that a storage gap covers the entire DCSS address range. See “Defining the guest storage with storage gaps.”
- OR-
- Enable Linux to handle real memory addresses that are beyond the guest storage (considered as real memory by Linux) to cover a DCSS that is located above the guest storage. See “Extending the Linux address range” on page 48.

Defining the guest storage with storage gaps

From a CMS session, use the DEF STORE command to define your guest storage as discontinuous storage extents. Ensure that the storage gap between the extents covers your entire DCSS. Issue a command of this form:

```
DEF STOR CONFIG 0.<storage_gap_begin> <storage_gap_end>.<storage above gap>
```

where:

<storage_gap_begin> is the lower limit of the storage gap. The lower limit must be at least 64 MB and at or below the lower limit of the DCSS.

<storage_gap_end> is the upper limit of the storage gap. The upper limit must be above the upper limit of the DCSS.

<storage above gap> is the amount of storage above the storage gap. The total guest storage is *<storage_gap_begin>* + *<storage above gap>*.

All values can be suffixed with M to provide the values in MB. Refer to the *z/VM: CP Commands and Utilities Reference* guide for more information on the DEF STORE command.

Example: Use the following command to create a DCSS that starts at 144 MB and ends at 152 MB accessible to a z/VM guest with 256 MB guest storage:

```
DEF STORE CONFIG 0.140M 160M.116M
```

The storage gap in the example ranges from 140 MB to 160 MB and thus covers the entire DCSS range. The total guest storage is 140 MB + 116 MB = 256 MB.

Extending the Linux address range

If your guest storage is sufficiently low, your entire DCSS address range might be above the guest storage. You can then modify the */etc/zipl.conf* file to make the DCSS accessible to the Linux guest. Perform these steps to extend the Linux address range:

1. Add the following command to the **[ipl]** section of the *parameters* line in the */etc/zipl.conf* file:

```
mem=<address>
```

where *<address>* is an address at or above the upper limit of the DCSS. Upper limits can be in KB or MB. **Important:** Be sure to include the **K** suffix or the **M** suffix after the address, as in the example that follows.

Example: To reserve storage for a DCSS that starts at 144 MB and ends at 152 MB accessible to a z/VM guest with 128 MB guest storage:

```
mem=160M
```

Note: For installations that use SUSE Linux Enterprise Server 10 for zSeries or for installations that use Red Hat Enterprise Linux 5, add two megabytes to the value specified for the *mem=<address>* statement.

2. Then, issue this command:

```
zipl
```

This command saves the configuration changes you just made.

CAUTION:

Be extremely careful when editing the */etc/zipl.conf* file. Any typographical errors made while editing the */etc/zipl.conf* file will disable your Linux system, rendering it unbootable.

3. Re-IPL your Linux guest.
4. To verify that the configuration change was successful, issue this command:

```
cat /proc/cmdline
```

This command displays your current configuration settings. Review this file to ensure that the *mem=* parameter was updated properly.

Step 6. Loading the DCSS device driver

Perform these steps to load and configure the DCSS device driver if the DCSS block device support has been compiled as a separate module.

1. Log in to the Linux guest as root.
2. Load the DCSS device driver with the *modprobe* command.

To load the DCSS module to the Linux guest, issue this command:

```
modprobe dcssblk
```

If the agent is running as a user other than root, change the ownership of the DCSS device to that user to allow the agent access to the DCSS. For example, if the agent is running as 'itmuser', issue this command:

```
chown itmuser /dev/dcssblk
```

Step 7. Adding the PERFOUT DCSS to your Linux guest

To add the PERFOUT DCSS to the Linux guest, run these commands on the Linux guest:

```
echo perfout > /sys/devices/dcssblk/add
```

Note: The *echo* command must be run as root on the Linux guest.

A message displays in the log that is located on */var/log/messages*. The message shows the range of addresses and the size of the PERFOUT DCSS. The message also states that the DCSS is now attached to this Linux guest.

Note: Only the Linux system where the OMEGAMON XE on z/VM and Linux monitoring agent is running must attach to the PERFOUT DCSS. Additionally, if the monitoring agent was installed as non-root, you must change the owner of this DCSS device driver to a non-root user; specifically, to the name of the user ID in the Linux guest under which it is running. This enables the monitoring agent to access the PERFOUT DCSS to retrieve the data for display on the Tivoli Enterprise Portal interface.

Step 8. Loading the PERFOUT DCSS at startup time

To configure the Linux guest so that it automatically loads the PERFOUT DCSS at startup time, edit the */etc/rc.d/boot.local* file to add these two commands:

```
modprobe dcssblk  
echo PERFOUT > /sys/devices/dcssblk/add
```

Save the *boot.local* file. The next time you reboot the system, you can verify whether or not the PERFOUT DCSS loaded automatically at startup time.

DCSS naming scheme

When you load a DCSS device driver, a major number is automatically allocated for it. A different major number can be used when the device driver is reloaded, for example when Linux is rebooted. To find out which major number is used for the DCSS, issue this command:

```
ls -l /dev/dcss*
```

When you add a DCSS, a minor number is assigned to it. Unless you use dynamically created device nodes as provided by *udev*, you might need to know the minor device number that has been assigned to the DCSS .

The standard device names are of the form *dcssblk<n>*, where *<n>* is the corresponding minor number. The first DCSS device that is added is assigned the name *dcssblk0*, the second *dcssblk1*, and so on. When a DCSS is removed, its device name and corresponding minor number are free and can be reassigned. A DCSS that is added always receives the lowest free minor number.

For detailed information on any of the procedures and commands used in these sections, refer to the *Device Drivers, Features and Commands* guide.

Step 9. Enabling the collection of Linux data

The **KVLUUser AppIDData** attribute group displays information about Linux memory usage, workload activity, and network utilization. Additionally, the **AppIDData** product-provided workspace displays Linux workload data using this attribute group. See the *IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide* for details on this attribute group and on this workspace.

Data collection for this attribute group must be enabled at each Linux guest system that is to be monitored. Enabling data collection for this attribute group defines the workgroup as a Linux workgroup. This is important because the links are displayed on a table that queries the systems in the Linux workgroup.

You must also enable the CP Monitor APPLDATA domain on the z/VM system where the Performance Toolkit is running. See “Enabling the CP Monitor domains” on page 40.

Additionally, before you can start the collection of data for this attribute group, you must update the directory entry for the Linux guest to include the APPLMON option. Use the OPTION operand of the DIRMAINT command to update the directory entry. Refer to the *z/VM V5R2.0 Directory Maintenance Facility Commands Reference* publication for details.

Starting data collection for User ApplData at the Linux guest

To start data collection for this attribute group, type the following commands on the Linux guest system:

```
modprobe appldata_os
modprobe appldata_mem
modprobe appldata_net_sum
echo 1 > /proc/sys/appldata/os
echo 1 > /proc/sys/appldata/mem
echo 1 > /proc/sys/appldata/net_sum
echo 10000 > /proc/sys/appldata/interval
echo 1 > /proc/sys/appldata/timer
```

Note: If you are running Red Hat Enterprise Linux v.4, you may not need to load the appldata modules by issuing the *modprobe* commands since the modules may already be built into the kernel. See “Required software” on page 21 for details on the supported platforms for this monitoring agent. If you are running on another supported platform, see “Determining if appldata modules are built into the Linux kernel” on page 51 for additional information.

Specify the command *echo <milliseconds> > /proc/sys/appldata/interval* in milliseconds of CPU time (not wall clock time). Additionally, these commands require root authority.

Important: You *must* issue, at a minimum, the *echo 1 > /proc/sys/appldata/os* command and the *echo <milliseconds> > /proc/sys/appldata/interval* command for all Linux guests that you want to monitor.

The interval set at the Linux guest for application monitor data collection determines how often new data are provided to the CP Monitor. You may want to set the interval at the Linux guest at a rate different than the CP Monitor sample interval. For example, setting the interval at the Linux guest to once every thirty minutes may be sufficient for your needs. It depends on your environment, and on how closely you want to monitor the Linux systems. Setting each interval independently provides greater control of data collection overhead.

The default setting at the Linux guest is 10,000 milliseconds (once every ten seconds) of CPU time (not wall clock time). This setting should offer an adequate compromise between performance and current data.

As an example, the following command sets the interval at the Linux guest to 20 seconds:

```
echo 20000 > /proc/sys/appldata/interval
```

The format of the command is as follows:

```
echo <milliseconds> > /proc/sys/appldata/interval
```

where *<milliseconds>* is the number of milliseconds of CPU time (not wall clock time) that elapse before new data are passed to the CP Monitor.

To determine the CP Monitor interval, issue a QUERY MONITOR INT command on z/VM. This command requests a display of the interval value currently in effect for sample monitoring on z/VM.

Determining if appldata modules are built into the Linux kernel

For supported platforms other than Red Hat Enterprise Linux v.4, to determine whether the appldata modules are built into the kernel, issue the following command as root:

To determine whether the appldata modules are built into the kernel, issue the following command as root:
`modprobe -l appldata*`

If output does not display, then the appldata modules are either missing or they are built into the kernel.

Issue the following command to ensure that the appldata modules are not missing:

```
ls /proc/sys/appldata
```

If the modules are built into the kernel, the files used to interface with the modules should be listed in the output. In particular, you should see listed the `appldata_os`, `appldata_mem`, and `appldata_net_sum` files. If these files are not listed, the appldata modules are missing from your installation. Contact your Linux vendor for assistance.

You must load the modules with the `modprobe` commands if the output of the `modprobe -l appldata*` command is similar to the following text:

```
linux-z2xg:~# modprobe -l appldata*
/lib/modules/2.6.16.21-0-8-default/kernel/arch/s390/appldata/appldata_os.ko
/lib/modules/2.6.16.21-0-8-default/kernel/arch/s390/appldata/appldata_mem.ko
/lib/modules/2.6.16.21-0-8-default/kernel/arch/s390/appldata/appldata_net_sum.ko
```

This output indicates that the appldata modules are not built into the kernel of your Linux installation, and you must issue the `modprobe` commands. Note that the kernel version (for example, 2.6.16.21-0.8) may differ from system to system.

See the *IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide* or the help system for descriptions of the KVLUserAppIData attribute group.

Enable the collection of Linux data automatically at startup time

Instead of having to manually enable the collection of Linux data each time you use the monitoring agent, you might want to enable it automatically at startup time.

To automatically enable the collection at startup time of the **KVLUser AppIData** attribute group that generates Linux data, add the following commands to your `/etc/rc.d/boot.local` file:

```
modprobe appldata_os
modprobe appldata_mem
modprobe appldata_net_sum
echo 1 > /proc/sys/appldata/os
echo 1 > /proc/sys/appldata/mem
echo 1 > /proc/sys/appldata/net_sum
echo 10000 > /proc/sys/appldata/interval
echo 1 > /proc/sys/appldata/timer
```

Save the `boot.local` file. The next time you reboot the system, review the **AppIData** workspace to verify that data are being collected.

Note: If you are running Red Hat Enterprise Linux v.4, you do not need to issue the *modprobe appldata ** commands prior to issuing the *echo ** commands. For all other supported platforms, you must issue the *modprobe appldata** commands. See “Required software” on page 21 for details on the supported platforms for this monitoring agent.

Important: You *must* issue, at a minimum, the *echo 1 > /proc/sys/appldata/os* command and the *echo <milliseconds> > /proc/sys/appldata/interval* command for all Linux guests that you want to monitor.

To access the **AppIData** workspace, from the Navigator, select the **Workload** workspace, select **Workspace**, and select **AppIData**. See the *IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide* for details on this attribute group and on this workspace.

Step 10. Enabling dynamic workspace linking

If you would like to link from the OMEGAMON XE on z/VM and Linux monitoring agent workspaces to the Tivoli Monitoring Agent for Linux OS workspaces, you must add an environment variable to the Tivoli Monitoring Agent for Linux OS configuration file. This step assumes that the Tivoli Monitoring Agent for Linux OS is installed and configured.

Important: If the Tivoli Monitoring Agent for Linux OS is started prior to adding the environment variable to the configuration file, a duplicate branch for the Linux OS agent may display in Tivoli Enterprise Portal. The original node will be dimmed in the Navigator. To clear the original node, right-click the Navigator item and click **Clear offline entry**. The original node will be deleted.

To add the environment variable that enables dynamic workspace linking between these two monitoring agents, perform the following steps:

1. Modify the **lz.ini** file. This file is located in the <ITM_Home>/config directory, where: <ITM_Home> is the directory where you installed IBM Tivoli Monitoring.
2. Add the statement **KLZ_SETLPARVMID=Y** to the **lz.ini** file.

Note: If you are running IBM Tivoli Monitoring, V6.2.0 and above, you also need to comment out the **CTIRA.HOSTNAME** variable.

3. Restart the Tivoli Monitoring Agent for Linux OS.

The agent will register with the Tivoli Enterprise Monitoring Server in the format:

```
LPAR.VM00ID:LZ
```

Note: Dynamic workspace linking between the OMEGAMON XE on z/VM and Linux monitoring agent and the Tivoli Monitoring Agent on Linux OS is not supported when the Linux system defined for the Linux OS agent is running as a guest under a second-level z/VM system.

Additionally, you must enable data collection for the KVL User Appldata attribute group at each Linux guest system that is to be monitored. Enabling data collection for this attribute group defines the workgroup as a Linux workgroup. This is important because the links are displayed on a table that queries the systems in the Linux workgroup.

You must also enable the CP Monitor APPLDATA domain on the z/VM system where the Performance Toolkit is running. For complete details on enabling the collection of Linux data, see “Step 9. Enabling the collection of Linux data” on page 49.

Important: Your Tivoli Enterprise Portal user ID must be authorized to access the Tivoli Monitoring Agent for Linux OS. Otherwise, links to workspaces in the Tivoli Monitoring Agent for Linux OS will not be included in the list of linked workspaces.

Choose a workspace from the list to navigate to that workspace. By linking to the target workspace in context, you receive additional information that is related to the system, subsystem, or resource you are currently viewing. If you choose a workspace from the list and the target workspace is not available, you receive message KFWITM081E. Refer to the *IBM Tivoli OMEGAMON XE on z/VM and Linux Troubleshooting Guide* for more information. See also the user's guides for both monitoring agents for descriptions of the workspaces provided with each monitoring agent.

Step 11. Enabling Take Action commands (optional)

The Take Action feature is optional. Configuring your environment for Take Action is only required if you intend to issue CP commands, CMS commands, and REXX executables from the Tivoli Enterprise Portal interface to the z/VM operating system. The CMS guest defined for Take Action operates as a dedicated, unattended, disconnected virtual machine that waits to process messages containing Take Action commands sent from the monitoring agent.

You can use the Take Action view to send commands to the CMS guest, to select a predefined action, and to stop or to start a process at the system. The predefined actions available depend upon where the window was invoked.

Select the Take Action view from the Tivoli Enterprise Portal toolbar using the Take Action icon. See the Tivoli Enterprise Portal online help for additional details about this feature.

Note: When specifying a Take Action command in Tivoli Enterprise Portal be sure to prefix all OMEGAMON XE on z/VM and Linux commands with **VL**:

There are certain guidelines to keep in mind when creating Take Action commands. Some types of commands need to be excluded from running. You can exclude commands from being executed by specifying them KVL CONFIG file. See “Guidelines for issuing Take Action commands” on page 56.

To enable Take Action commands, several requirements must be met. These are outlined in the sections that follow.

Note: When you create a situation or edit a predefined situation, you can also issue commands using the **Action** tab of the Situation Editor. The command is sent to the system when the situation becomes true. The requirements that apply to Take Action commands apply equally to using the Action tab to issue commands that are executed on the z/VM operating system.

z/VM requirements for Take Action commands

The requirements that follow for Take Action commands pertain to the z/VM operating system:

- The Take Action Command Processor REXX executable file (KVL CMD EXEC) that ships with this agent must be installed. See the *Program Directory* for instructions on installing this executable file. See also the z/VM Installation section of “The IBM Tivoli OMEGAMON XE on z/VM and Linux product package” on page 26 for the contents of the tape that contains the Command Processor files.
- KVL CONFIG - Modify this file to contain a list of the allowed user IDs (AGENT_ID = the VM user ID of the Linux guest where this monitoring agent is running). This file also includes a list of the commands that are not allowed, for example, CMDS=LOGOFF.

You can also specify logging the output from the processed commands to a log file. You can specify the following settings:

- LOG_SIZE=<number of lines>. The log size range is 100 - 5000 lines. The default number of lines is 5000 lines.
- LOG_COUNT=<number>. The log cycles through as many times as you specify in this setting. For example, KVL 1, KVL 2, and so on. A maximum of five active log files is retained.
- LOG_RESP=Y or N. You specify whether or not the results of the processed commands are written to the log file. The default is **N** - the output will not be written to the log file.**Important:** Turning

results logging on will result in significantly larger log files. Modify the settings to meet the needs of your environment. The log file generated is called KVL <number>, where number is the number of times the log file has been generated.

- You must create and configure a Conversational Monitor System (CMS) guest system. The CMS guest system must reside on the same host on which the monitoring agent resides. The user ID must be set up as an unattended, disconnected service machine. The user ID should have the privileges it will require for the types of commands you want to issue by means of the Take Action commands.

Setting the environment variables during monitoring agent installation

During the installation of this monitoring agent, you are prompted to enter values for the following environment variables:

- KVL_CMDUSERID - Set this environment variable with the z/VM CMS guest user ID of the system where the Command Processor was installed. If the CMS guest user ID is not specified, the Command Processor will be disabled. Take Action events initiated by situations or submitted by you using the Take Action feature will not be processed.
- KVL_MSGTYPE - To describe the type of message being used to process the commands, set this environment variable to a value of either

msg

or

smsg

When the value is set to **smsg**, the monitoring agent uses the CP SMSG command to send Take Action commands to the Command Processor virtual machine.

When the value is set to **msg**, the monitoring agent uses the CP MSG command to send Take Action commands to the Command Processor virtual machine. If a value is not specified, the Take Action feature is disabled.

The results of the Take Action commands can be viewed in the KVL log file, located on the guest system where the monitoring agent is installed.

Linux guest requirements for Take Action commands

The following requirements for Take Action commands pertain to the Linux guest hosting the OMEGAMON XE on z/VM and Linux monitoring agent:

- SUSE Linux Enterprise Server 9 for zSeries, with Service Pack 3 or later must be installed on the Linux on zSeries image where the z/VM agent is running. SUSE Linux Enterprise Server 10 for zSeries is also supported, as is Red Hat Enterprise Linux v.4 Update 5. Red Hat Enterprise Linux 5 is also supported.
- The z/VM CP interface device driver (vmcp) must be available on the Linux guest on which the monitoring agent is running. The vmcp driver allows the monitoring agent to issue CP commands, thus ensuring that Take Action commands are sent to the designated z/VM CMS guest user for processing.
- sudo, the superuser do utility for Linux-based systems, is required for the monitoring agent to run with temporary root authority to send Take Action commands to the Command Processor virtual machine. The sudo utility must be available and configured on the Linux guest on which the monitoring agent is running. This utility allows a non-root Linux guest user to run vmcp without having root authority. The sudo utility is generally pre-installed and available as part of the Linux operating system. Any working version is sufficient.

Note: If you are running as a root user, you must perform Step 1 of the procedures that follow for SUSE Linux Enterprise Server 10 for zSeries and for Red Hat Enterprise Linux 5.

Running under Linux as a non-root user

If you are running under Linux as a non-root user, please review the man pages for sudo, visudo, and for sudoers to make sure you understand their usage. In the following sections, we assume that the non-root user is named **tivoli**. If your non-root user is named otherwise, please edit the commands accordingly.

Note: If you are running as a root user, you must perform Step 1 of the procedures that follow for either SUSE Linux Enterprise Server 10 for zSeries or for Red Hat Enterprise Linux 5.

On SUSE Linux Enterprise Server 9 for zSeries: On the Linux system where you installed this monitoring agent, modify the file called `/etc/sudoers`. You do this by adding an entry specifying the command that this user is allowed to run as root. Use only `visudo` to edit the `/etc/sudoers` file by adding the following command:

```
tivoli ALL=NOPASSWD:/usr/bin/vmcp
```

On SUSE Linux Enterprise Server 10 for zSeries: Perform the following steps on the Linux system where you installed this monitoring agent:

1. Modify the `vl.ini` file. This file is located in the `ITMinstall_dir/config` directory, where: `ITMinstall_dir` is the directory where you installed IBM Tivoli Monitoring. Locate the `PATH=` environment variable for the non-root user and add the following statement:

```
:/sbin
```

Note: All SUSE Linux Enterprise Server 10 for zSeries installations must perform Step 1, as it applies to root users as well as to non-root users.

2. Modify the file called `/etc/sudoers` by adding an entry for the Linux user specifying the command that this user is allowed to run as root. Use only `visudo` to edit the `/etc/sudoers` file by adding the following command:

```
tivoli ALL=NOPASSWD:/sbin/vmcp
```

3. Restart the monitoring agent on the affected systems by issuing the following command:

```
./itmcmd agent start pc
```

where `pc` is the product code for the agent that you want to start. The product code for this monitoring agent is `vl`. The product code for the Tivoli Monitoring Agent for Linux OS is `lz`. For a list of product codes, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

On Red Hat Enterprise Linux 5: Perform the following steps on the Linux system where you installed this monitoring agent:

1. Modify the `vl.ini` file. This file is located in the `ITMinstall_dir/config` directory, where: `ITMinstall_dir` is the directory where you installed IBM Tivoli Monitoring. Locate the `PATH=` environment variable for the non-root user and add the following statement:

```
:/sbin
```

Note: All Red Hat Enterprise Linux 5 installations must perform Step 1, as it applies to root users as well as to non-root users.

2. Modify the file called `/etc/sudoers` by adding an entry for the Linux user specifying the command that this Linux user is allowed to run as root. Use only `visudo` to edit the `/etc/sudoers` file by adding the following command:

```
tivoli ALL=NOPASSWD:/sbin/vmcp
```

3. Restart the monitoring agent on the affected systems by issuing the following command:

```
./itmcmd agent start pc
```

where `pc` is the product code for the agent that you want to start. The product code for this monitoring agent is `vl`. The product code for the Tivoli Monitoring Agent for Linux OS is `lz`. For a list of product codes, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

See “Step 11. Enabling Take Action commands (optional)” on page 53 for details on enabling Take Action commands.

Loading the vmcp device driver

To load the vmcp device driver, issue this command at the Linux guest:

```
modprobe vmcp
```

Verifying that sudo is available

To verify that sudo is available, issue this command on the Linux guest where the monitoring agent is installed:

```
sudo vmcp q userid
```

If sudo is active and loaded, this command sends the **q userid** command to the hosting virtual machine, which queries the user ID for the guest.

Verifying that Take Action commands are being issued

To verify that the monitoring agent is sending commands to the Command Processor, issue this command at the Linux guest:

```
sudo vmcp smsg <userID> <TEP userID> cmd=<name of the command>
```

Where:

- userID = the user ID of the CMS guest where the Command Processor is running.
- TEP user ID = the user ID of the person authorized to issue Take Action commands in the Tivoli Enterprise Portal. Review the Tivoli Enterprise Portal online help to find out how to assign permissions.
- name of the command = the identifier for the command you want the Command Processor to execute.

Note: The format of this command varies depending on whether you are using `msg` or `msg`. See “Setting the environment variables during monitoring agent installation” on page 54 for details.

Guidelines for issuing Take Action commands

With this monitoring agent, you can issue any of the CP and CMS commands as Take Action commands. Additionally, you can also invoke your own REXX executables as Take Action commands. All REXX executable files must be pre-installed on the CMS guest system prior to being submitted as Take Action commands.

This section lists some guidelines to keep in mind when creating Take Action commands. It describes the types of commands that need to be excluded from running, and includes the sample commands to be denied execution that appear in the KVL CONFIG file.

Some commands when issued at the monitoring agent could have unpredictable results, and therefore are not supported.

The types of commands that are not supported by this monitoring agent are the following:

- Full-screen commands or executable files
- Long-running commands or executable files
- CMS immediate commands (for example, HB, HI, HO, HT, HX, RO, RT, SO, TE, and TS)

Note: An immediate command is not recognized by the Programmable Operator Facility if it is being used.

- Commands or executable files that cause a VM READ or CP READ

Note: Such a command will stop the operation of the Programmable Operator Facility if it is being used.

- Commands that alter or overlay CMS storage (for example, CP DEFINE STORAGE, CP IPL CMS, CP LOGOFF, CP RESET, CP DEFINE STORAGE)

- Commands that contain line-editing characters (pound sign (#), for example), as defined by the CP TERMINAL command

Note: Such characters are not recognized as line-editing characters by the Programmable Operator Facility, if it is being used.

The results of the Take Action commands are logged in the KVL log file. This log file is located in the same location where the monitoring agent is installed. See “z/VM requirements for Take Action commands” on page 53 for details on the log file.

For information on the CP commands and on the CMS commands, refer to the *z/VM: CP Commands and Utilities Reference* and to the *z/VM: CMS Commands and Utilities Reference* guides, respectively.

Commands to be excluded from execution

Some commands need to be excluded from running on the CMS guest system. Denied operating system commands are not issued as Take Action commands because they pose an obvious danger to the continued operation of the system. Denied commands are filtered out on the CMS guest system and are not invoked.

Modify the file called KVL CONFIG, that ships with this monitoring agent, to include the list of commands that you do not want to run as Take Action commands. The KVL CONFIG file contains sample commands that are prevented from running. You can add more commands to this file to prevent those commands from being issued. You can also remove any of the sample commands in this file. The following is a sample of an excluded command in the KVL CONFIG file:

```
CMDS=LOGOFF
```

The five sample commands that are prevented from being run as Take Action commands are the following:

- CP
- IPL
- LOGOFF
- RESET
- DEFINE STORAGE

You also use the KVL CONFIG file to add the user IDs of the Linux guests on which the monitoring agent is running. See the z/VM Installation section of “The IBM Tivoli OMEGAMON XE on z/VM and Linux product package” on page 26 for details on the KVL CONFIG file.

Step 12. Installing the Language Packs (optional)

If you want the OMEGAMON XE on z/VM and Linux monitoring agent workspaces, online help, and expert advice to be displayed in a language other than English, you can install language support for the monitoring agent on all the workstations where a Tivoli Enterprise Portal Desktop client is located and where the Tivoli Enterprise Portal Server is installed. Additionally, for best results, after you apply an IBM Tivoli Monitoring fix pack and reconfigure any of the base components, re-install the base IBM Tivoli Monitoring language packs and any agent language packs.

If IBM Tivoli Monitoring language support has not yet been installed, you must install it before installing monitoring agent language support. Language support is available only on the platforms that were supported for IBM Tivoli Monitoring v6.x GA. See the *IBM Tivoli Monitoring Installation and Setup Guide* for instructions.

Important: If you already have IBM Tivoli Monitoring language support installed, you will need to upgrade your monitoring agent language support files to ensure you are running with the same level of language support as Tivoli OMEGAMON XE on z/VM and Linux, version 4.2.0. To run with the same level of

language support as Tivoli OMEGAMON XE on z/VM and Linux, version 4.2.0, install the latest IBM Tivoli Monitoring language support before installing the monitoring agent language support.

Use the following steps to install the monitoring agent language pack:

1. Insert the *IBM Tivoli OMEGAMON XE on z/VM and Linux Language Pack* CD into the CD-ROM drive of the computer where the Tivoli Enterprise Portal Server and the Tivoli Enterprise Portal desktop client are located. For Linux/UNIX, mount the CD, if necessary.
2. In the directory where you extracted the language pack installation image, launch the installation program as follows:
 - On Windows: Setup should begin automatically. If setup does not begin automatically, go to the Windows directory on your CD-ROM drive and run the `lpinstaller.exe` command.
 - On Linux/UNIX: Run the `./lpinstaller.sh` command.
3. Select the language to be used during the installation, and click **OK**.
4. Read the text that welcomes you to the installation, and click **Next** to continue.
5. If you are updating a language pack or adding a new language pack, click the **Add/Update** icon and click **Next**.
6. Select the directory in which to install the National Language pack (NLSPackage) files. The files must be installed in the directory where the Tivoli Enterprise Portal Server and the Tivoli Enterprise Portal desktop client are located. If the desktop client is on a different machine, install the language files there as well. Click **Next**.
7. Select the monitoring agent for which you would like to install a language pack and click **Next**.
8. Select one or more languages from the list of supported languages, then click **Next**.
9. Review the installation summary and click **Next**. The language support files are installed, and a message instructs you to restart the Tivoli Enterprise Portal Server and the Eclipse Help Server.
10. Click **Done** to exit the installer.
11. Stop and restart the following components:
 - Tivoli Enterprise Portal Server
 - Eclipse Help Server
 - Tivoli Enterprise Portal desktop or browser client

For instructions on specifying the language to be displayed for users, see *IBM Tivoli Monitoring Administrator's Guide*.

Starting the monitoring agent

This section assumes that you have installed and configured the monitoring agent. See the *IBM Tivoli Monitoring Installation and Setup Guide* for details.

Use the **itmcmd agent** command at the command line to start and stop the monitoring agent. You can start or stop one agent, all agents, or multiple agents. You can also start the portal server and portal desktop client using this command.

You must run the **itmcmd agent** command on the architecture for which the agent is installed.

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. The product code for this monitoring agent is **vl**. The product code for the IBM Tivoli Monitoring Agent for Linux OS is **lz**. For a complete list of product codes, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

Stopping the monitoring agent

To stop all monitoring agents, run the following command:

```
./itmcmd agent stop all
```

To stop specific agents, run the following command:

```
./itmcmd agent stop pc pc pc
```

where *pc* is the product code for the agent that you want to start. The product code for this monitoring agent is **vl**. The product code for the IBM Tivoli Monitoring Agent for Linux OS is **lz**. For a complete list of product codes, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

Chapter 6. Defining user IDs and security

This chapter provides information about how to define the list of user IDs that the Command Processor uses for Take Action commands. This chapter also describes how user IDs and security are implemented in IBM Tivoli Monitoring.

Defining the list of user IDs for the Command Processor

The Command Processor associated with the Take Action feature requires a list of user IDs for the Linux guests that are allowed to issue Take Action commands. You assign user IDs for the Linux guests using the KVL CONFIG file that is part of the product package. The Take Action feature is optional. Configuring your environment for Take Action is only required if you intend to issue CP commands, CMS commands, and REXX executables from the Tivoli Enterprise Portal interface to the z/VM operating system.

Use the sample KVL CONFIG file that is provided with this monitoring agent to define the list of allowed user IDs for Linux guests. This file contains a parameter called AGENT_ID, where AGENT_ID = the user ID of the Linux virtual machine where this monitoring agent is running.

The KVL CONFIG file also includes a sample list of the commands that are to be excluded from running. For example, CMDS=LOGOFF. Modify this file to include the list of commands that you want to exclude from running as Take Action commands.

See the *Program Directory* for information on the KVL CONFIG file.

Defining user IDs in IBM Tivoli Monitoring

During the planning phase, consider which users will require access to the Tivoli Enterprise Portal and which features, applications and views you will permit users to access.

Each user needs a user ID to access performance data through the Tivoli Enterprise Portal. With the user administration function provided by Tivoli Enterprise Portal Server, you can control access to features, applications, and views by defining user IDs and assigning access through the Tivoli Enterprise Portal user interface. Initially, the Tivoli Enterprise Portal Server has only one valid user ID: **sysadmin**. This user ID allows the administrator to log on and create other user IDs.

Password protection is provided by the hub Tivoli Enterprise Monitoring Server. During user logon to the Tivoli Enterprise Portal, the user ID is validated by the Tivoli Enterprise Portal Server and, if security is enabled at the hub Tivoli Enterprise Monitoring Server, the user ID and password are validated by the hub Tivoli Enterprise Monitoring Server. The hub Tivoli Enterprise Monitoring Server validates the user ID and password with the local operating system. Monitoring agents such as OMEGAMON XE on z/VM and Linux do not provide additional security mechanisms.

The first time you configure the Tivoli Enterprise Monitoring Server, configure it with security turned off. Perform the following steps to enable security:

- Configure all products and verify that they are operating properly.
- If you choose a third-party security package, verify that it is installed and configured properly for your site.
- Create user IDs and passwords on the system running the hub Tivoli Enterprise Monitoring Server. Authorize the users to the resources they will access.
- Change your hub Tivoli Enterprise Monitoring Server configuration to enable security and specify which security product will be used.

When you are running on a Windows server, the hub Tivoli Enterprise Monitoring Server uses the Windows operating system security APIs that validate the user ID and password in the following sequence:

1. Local workstation
2. Domain controller if the local workstation is part of a domain
3. Any domain controller that has established a trusted relationship with the local domain

Tivoli Enterprise Portal users do not require Windows administrative authority on the Windows system running a hub Tivoli Enterprise Monitoring Server. The minimum requirement is that the user ID must be granted the *Log on locally* user rights policy on the hub Tivoli Enterprise Monitoring Server. Refer to your Windows documentation to learn how to configure security on a Windows system and to the configuring user security topic in *IBM Tivoli Monitoring Installation and Setup Guide*. See the *IBM Tivoli Monitoring Administrator's Guide* for detailed information about creating and maintaining users.

A hub Tivoli Enterprise Monitoring Server running on z/OS validates user IDs and passwords using either the security feature Network Access Method (NAM) or through a supported system authorization facility (SAF) product. For more information about security on a Tivoli Enterprise Monitoring Server running on z/OS, refer to *Configuring IBM Tivoli Enterprise Monitoring Server on z/OS*.

Part 3. Completing your configuration

The chapters in this section describe the process of completing the configuration of your monitoring agent.

- Chapter 7, “Performance and storage considerations,” on page 65 identifies options to consider while completing the configuration and deployment of this product to meet the needs of your enterprise.
- Chapter 8, “Serviceability,” on page 73 describes the product features, tools, and documentation that relate to troubleshooting, problem determination, or problem source identification.

Chapter 7. Performance and storage considerations

This chapter identifies options to consider while completing the configuration and deployment of this product to meet the needs of your enterprise. This chapter also provides information to consider when creating and modifying situations and workspaces.

Collecting, processing and storing data consumes system and network resources. There are several variables that contribute to the amount of CPU, memory and disk space used by OMEGAMON XE on z/VM and Linux and the IBM Tivoli Monitoring components. The amount of CPU, memory and disk space used on each monitored system will depend on the number of resources being monitored, how often performance data are collected and whether or not you choose to store historical data.

When a performance bottleneck is suspected, that is, when interactive system response time is degrading or when job turnaround times become too long, the reason for this unsatisfactory system performance has to be determined so actions can be defined for improving performance. Information in the following publications will give you some general guidance on how to determine the causes of unsatisfactory system performance.

The following manuals contain z/VM operating system specific tuning hints:

- *z/VM: CP Planning and Administration*
- *z/VM: Running Guest Operating Systems*

Additionally, performance information for the OMEGAMON XE zSeries products is available at <http://www.ibm.com/software/tivoli/features/ccr2/info.html>.

Understanding how data are collected

Using data collection, you can monitor the performance of your network and system resources in an effort to resolve potential problems before they affect the end user. The OMEGAMON XE on z/VM and Linux monitoring agent is designed with performance in mind, providing an efficient mechanism for collecting large amounts of network performance data.

By default, the OMEGAMON XE on z/VM and Linux monitoring agent is configured to display various types of data, such as the following:

- CPU
- Channel load
- Control unit cache
- DASD
- DASD cache
- FICON channel load
- HiperSockets
- LPAR
- LPAR channel load
- Linux workload or Linux machine
- Minidisk cache utilization
- Network
- Paging and spooling
- Real storage
- Specialty engine
- Spin lock

- TCP/IP
- Virtual disk
- Virtual switch
- Workload

The data are collected from a defined z/VM PERFOUT DCSS by the Performance Toolkit. Data collection is linked to the monitoring sample interval set at the CP Monitor.

Determining which systems to monitor

All production LPARS should be monitored since they represent the core of your business enterprise. You might choose to collect data less frequently on some systems, especially non-production LPARs, to minimize the cost of monitoring your networks.

Understanding historical data

Historical data collection is an optional feature that is enabled and configured using the Tivoli Enterprise Portal. The IBM Tivoli Monitoring provides the following types of historical data collection:

- Short-term historical data are stored in files on distributed systems or in the persistent datastore on z/OS systems.
Short-term historical data usually refers to data that is stored for 24 hours or less. Beyond 24 hours, the oldest data samples are deleted as new ones arrive, or if you have a data warehouse, the data are rolled off before being deleted.
- Long-term historical data are stored in the Tivoli Data Warehouse, which resides on a Windows computer. The long-term history database can retain data collected by OMEGAMON XE on z/VM and Linux monitoring agents for as long as you like (days, weeks, months or years).

Short-term historical data are best used for analysis during problem determination.

Long-term historical data are better used for trend analysis and to determine workload balance. Long-term historical data collection requires installation of a relational database. The Tivoli Data Warehouse uses a DB2, Oracle, or Microsoft SQL Server database to store historical data collected across your environment. You can generate warehouse reports for short term and long-term data through the Tivoli Enterprise Portal. You can also use third-party warehouse reporting software, such as Crystal Reports or Brio, to generate long-term data reports. Warehouse reports provide information about the availability and performance of your monitoring environment over a period of time.

The Tivoli Data Warehouse uses the Warehouse Proxy agent to move data from monitoring agents or the monitoring server to the data warehouse database. The Warehouse Proxy is an ODBC export server for warehousing historical data. It is a special agent that uses an ODBC connection to transfer historical data collected from agents to a database. You can then analyze this data using the workspaces in the Tivoli Enterprise Portal or any third-party software. Short-term historical data collection must be enabled and configured if you want to perform long-term historical data collection.

The Summarization and Pruning Agent is a mechanism for managing data in the Tivoli Data Warehouse. The data in the Tivoli Data Warehouse is a historical record of activity and conditions in your enterprise. Summarization of the data is the process of aggregating your historical data into time-based categories. For example, hourly, daily, weekly, and so on. Summarizing data allows you to perform historical analysis of the data over time. Pruning of the data keeps the database to manageable size and thus improves performance. Pruning of the database should be performed at regular intervals.

Important: You can run only one Summarization and Pruning agent even if you have multiple monitoring servers that are sharing a single Tivoli Data Warehouse database. Running multiple Summarization and Pruning agents causes conflicts because the multiple instances attempt to prune the data in the tables simultaneously. The negative impact is that the configuration settings for the summarization and pruning

periods need to be set only in one monitoring server - that monitoring server controls how the data are summarized and pruned for all monitoring servers. See the *IBM Tivoli Monitoring Administrator's Guide* for details on Summarization and Pruning.

After historical data collection is enabled, an icon is displayed in qualifying views in Tivoli Enterprise Portal workspaces. Allow time for historical data to be stored, to produce meaningful reports. You can click on this icon to extend any existing Tivoli Enterprise Portal view (also called a report) to include historical data. Tivoli Enterprise Portal reports automatically pull data from short-term and long-term history, based upon the time period you specify for the report.

The collection interval for historical data can be configured to be different than the collection interval for real-time data. To avoid processing overhead and decrease storage consumption, historical data collection is typically performed less frequently than real-time data collection. You can configure a short-term historical data collection interval of 5, 15, 30 or 60 minutes.

Writing the data to long-term history can be configured for 24 hours, 1 hour or OFF. If you configure long-term history, use a warehousing interval of 1 hour to avoid transferring 24 hours worth of historical data at one time. This shorter interval will reduce the duration of CPU usage associated with writing data to the warehouse by spreading the writing across 24 periods.

Determining which types of historical data to collect

When deciding which types of data to store in short-term and long-term history and how long to store it you need to recognize that data collection consumes CPU cycles and disk space. Writing data to short-term history is cost effective and typically much less costly than writing to long-term history.

Short-term historical data are written to disk, performed either at the monitoring agent or at the monitoring server, consuming CPU cycles. Additional CPU cycles are used when the Warehouse Proxy extracts data from short-term history and transfers it to the Data Warehouse.

Depending on your needs, you may configure historical data collection for only a subset of attribute tables. This is an effective means for limiting storage consumption, particularly if you choose not to perform historical data collection for high volume attribute tables such as TCP connections or attribute tables with many bytes per row (many attributes), such as System attributes. Do not collect data you will not use in historical reports.

Use the tables in "Disk capacity planning for historical data" on page 68 to calculate storage consumption based on real-time data collection. The information you gather for these tables also provide the basis for calculating the storage requirements for historical collection.

You can use this information as a basis for choosing which attribute tables to enable for historical collection. You can select individual attribute tables for historical collection, including specifying different historical collection intervals and warehouse intervals.

You may elect to not enable collection of all attribute tables. This choice affects which workspaces can draw historical reports, but the data are there if needed for reporting.

By default historical reports retrieve up to 24 hours of data from short-term history.

Because historical data accumulates, you must determine how long you want to keep the data. You can also write short-term history to flat files, for backup, or for analysis in a statistical or graphing package.

Long-term history, in a SQL database, does not automatically prune old records. You need to determine how much data to retain, and either use the database manager tools to manually delete old records or schedule a script to run automatically to delete old records.

For detailed instructions about setting up historical data collection, refer to the *IBM Tivoli Monitoring Installation and Setup Guide*. Additional agent configuration information about the warehouse proxy is found in the *IBM Tivoli Monitoring Administrator's Guide*. For information about reporting, refer to *IBM Tivoli Monitoring Enterprise Portal User's Guide*.

Disk capacity planning for historical data

Disk capacity planning for a monitoring agent is a prediction of the amount of disk space to be consumed for each attribute group whose historical data are being collected. Required disk storage is an important factor to consider when you are defining data collection rules and your strategy for historical data collection. Calculate expected disk space consumption by multiplying the number of bytes per row by the expected number of rows, and then multiplying that product by the number of samples.

The following table shows the storage costs for the required and optional types of data. These tables are provided to inform you of the relative size of attribute tables. You may use this information to determine which types of data to monitor. The data shown in Table 6 is collected once every collection interval.

- *Attribute group* is the name for the type of data being monitored.
- *DB table name* is the table name as it would appear in the warehouse database, if the attribute group is configured to be written to the warehouse.
- *Row size in bytes* is an estimate of the record length for each row or instance written to the agent disk for historical data collection. This estimate can be used for agent disk space planning purposes.
- *Summarization row size in bytes* is an estimate of the record length for each row or instance written to the agent disk for summarized historical data. This estimate can be used for agent disk space planning purposes.
- *Expected number of rows* is a guideline that can be different for each attribute group, because it is the number of instances of data that the agent will return for a given attribute group, and depends upon the application environment that is being monitored.

The *IBM Tivoli Monitoring Installation and Setup Guide* contains formulas that can be used to estimate the amount of disk space used at the agent and in the warehouse database for historical data collection of an attribute group.

Table 6. Capacity planning for historical data

Attribute group	DB table name	History row size in bytes	Summarization row size in bytes (Daily, Hourly, Weekly, Monthly)	Total History and Summarization	Expected number of rows
KVLChannel	CHANNEL	504	1652	7112	1 row per z/VM channel
KVLCPU Device	VMCPDEV	164	394	1740	1 row per CP-owned device allocation
KVLControlUnit	CTLUNIT	702	2672	11390	1 row per DASD caching control unit
KVLDASDCache	DASDCACHE	1218	4412	18866	1 row per DASD cache
KVLDevice	VMDEV	303	998	4295	1 row per DASD device
KVLFChannel	FCHANNEL	403	1272	5491	1 row per FICON channel
KVLHiperSocket	VMHIPER	341	1024	4437	1 row per HiperSockets channel
KVLLChannel	LCHANNEL	508	1656	7132	1 row per LPAR channel
KVLMinidisk Cache	MDCACHE	920	3790	16080	1 row per system
KVLLPAR Info	LPARINFO	578	1819	7854	1 row per LPAR
KVLPKStat	PTKSTAT	160	192	928	Varies, depending upon your environment
KVLProcessor Data	PROCESSOR	441	1217	5309	1 row per processor engine
KVLSpinLock	SPINLOCK	403	1272	5491	1 row per lock name
KVLSystem	VMSYSTEM	700	4950	19300	1 row per system
KVLSystem2	VMSYSTEM2	326	1976	8230	1 row per system
KVLTCP/IP Srvr Data	TCPDATA	676	3120	13156	1 row per TCP/IP server
KVLTCP/IPUsrData	TCPUDATA	319	1212	5167	1 row per TCP/IP user

Table 6. Capacity planning for historical data (continued)

Attribute group	DB table name	History row size in bytes	Summarization row size in bytes (Daily, Hourly, Weekly, Monthly)	Total History and Summarization	Expected number of rows
KVLUser ApplData	VMLXAPPL	1374	5138	21926	1 row per Linux guest system (with APPLDATA enabled)
KVLUser Wait	VMWAIT	212	212	1060	1 row for every virtual machine guest logged into the system + 1 row to encompass the overall system
KVLUser Workload	VMWORK	531	2980	12451	1 row for every virtual machine guest logged into the system
KVLVDisk	VDISK	408	1487	6356	1 row per virtual disk
KVLVirtualSwitch	VSWITCH	851	3046	13035	1 row per virtual switch

Defining and running situations

Situations are used to identify monitored resources that meet certain performance criteria, raising an alert when the criteria is met. A situation definition includes a sampling frequency, a set of conditions, and a list of monitored systems. Each of these has implications on CPU and storage consumption. Also, the cumulative effect of all the active situations has implications on performance.

OMEGAMON XE on z/VM and Linux provides predefined situations to detect some of the most common problems. The majority of situations are disabled by default at installation, and will not start automatically until configured to do so. Refer to the *IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide* and to the agent online help for descriptions of the product-provided situations.

When planning for configuration and deployment of OMEGAMON XE on z/VM and Linux, evaluate all situations provided by the product. Determine which situations to autostart and eliminate situations that are not relevant to your monitoring strategy. If necessary, modify existing situations and create new situations to meet the needs of your enterprise.

For each situation that you choose to run, determine the importance and therefore the desired sampling frequency. Selecting several different intervals can increase the cost without increasing the benefit. Consider the attribute table that a situation queries. A good way to optimize the performance of situations is to enable the Tivoli Enterprise Portal to group situations. Do not set the sampling frequency to one that is shorter than the CP Monitor sampling interval. The additional samples return the same data until the CP Monitor provides a new sample.

For situations to be grouped, they must be active when the hub monitoring server starts, have the same sampling frequency, and test conditions on attributes in the same attribute table. Certainly you want to autostart all high-severity situations on a group with the same sampling frequency. You may also be able to gain efficiencies by auto-starting all situations on a group with the same sampling frequency, regardless of severity.

Verify that the conditions evaluated by each situation are appropriate for your environment. Check both the set of conditions and individual conditions. The predefined situations attempt to use the most efficient means to identify problems. In your environment there might be alternative conditions that identify the same problems but are less expensive to evaluate. Ensure that the values being checked are correct for your environment.

After you modify situations that are auto-started, stop and start the hub monitoring server. The process of combining situations occurs only during initialization of the hub monitoring server. Refer to the Tivoli Enterprise Portal online help for information relating to situations.

Designing workspaces

When you navigate to a workspace, one or more queries are processed by the IBM Tivoli Management Services components to display the requested workspace. Those same queries are processed again when you request a refresh or periodically in the cases where the workspace is configured to refresh automatically.

The workspaces and queries provided in the OMEGAMON XE on z/VM and Linux product have been designed with performance in mind. However, your environment and the resources you monitor might require customization of the product-provided workspaces and queries.

The following are tips to improve the performance for viewing workspaces. Refer to *IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide* to see descriptions of the workspaces provided with this monitoring agent.

The query assigned to a chart or table view requests data from a particular attribute table. The query runs when you open or refresh the workspace. The portal server sends the query to the hub monitoring server. The hub monitoring server distributes the query to the appropriate monitoring agent or agents and aggregates the resulting rows. The portal server retrieves the results and holds the entire result set in memory. The Tivoli Enterprise Portal retrieves one page of the results to display and holds both the current and previous page in memory.

You can dramatically reduce the amount of data retrieved by doing the following:

- Reducing the number of rows or attributes retrieved
- Applying the same query to multiple views in a workspace
- Adjusting the auto-refresh rate

Reducing the number of rows retrieved

One of the best ways to improve the performance of queries is to reduce the number of rows retrieved. Use the Query Editor to add filters that reduce the number of rows that are returned. You might want to change the existing filter values of a query or add filters to the query. For example, the LPAR workspace contains a table view that displays data for all logical partitions being monitored. You might be interested in only those logical partitions that are active. You could customize the query by adding a filter in the Query Editor for the LPAR Status attribute by selecting **Active**. Refer to the *IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide* to see descriptions of each workspace. Refer to the Tivoli Enterprise Portal online help for information on defining queries.

Do not confuse custom queries with view filters, which can also be invoked from the Tivoli Enterprise Portal Properties dialog. View filters have no effect on reducing the CPU and storage consumption by the monitoring agent and actually increase the Tivoli Enterprise Portal client CPU consumption.

View filters are applied by the client and affect only the current page. If more than one page is returned by the query, only a subset of the data is viewed on each page. Increasing the page size is an option available in Tivoli Enterprise Portal. This typically provides more filtered data on each page, but increases the client's memory requirements as now the two pages per query stored at the client are larger. It is more efficient to filter in the queries.

Reducing the number of attributes retrieved

Most product-provided queries return all attributes. There might be 50 attributes in an attribute table, yet you might want to view only 25 of them. Creating a custom query to retrieve only those 25 attributes reduces Tivoli Enterprise Portal Server and client processing and memory requirements.

For example, the OMEGAMON XE on z/VM and Linux System attribute table contains around 45 attributes. You can create a workspace that displays only those attributes for which you are most interested in seeing data. Refer to the Tivoli Enterprise Portal online help for more information about modifying or creating queries.

Applying the same query to multiple views in a workspace

Having multiple views in a workspace that retrieve data from different attribute tables is acceptable. However, if you have two views containing attributes that are available from the same attribute table, you must create one custom query for both views. By creating a single custom query, Tivoli Enterprise Portal will retrieve the data once for both views.

The objective is to use only one query for each attribute table used in a workspace. When a workspace is displayed, the entire results set for each query is stored on the Tivoli Enterprise Portal Server. The 100 rows (default page size) from each query currently being viewed and the previous page of any pane viewed are stored on the Tivoli Enterprise Portal client.

Adjusting the auto-refresh rate

You can choose an automatic refresh rate from every 60 seconds to once per hour. Each time the workspace is refreshed, the data are retrieved from the system where the monitoring agent is running. Retrieving data from the agent consumes CPU so it is important to specify a refresh rate that meets your monitoring needs while avoiding unnecessary performance overhead by the monitoring agent.

Chapter 8. Serviceability

Serviceability is defined as product features, tools, and documentation relating to troubleshooting, problem determination, or problem source identification. On a broader level, it can also encompass service offerings or processes that make the product more serviceable.

To make the Tivoli monitoring agents serviceable, log files containing messages and trace information are provided in a common fashion across all the monitoring agents and the components of IBM Tivoli Monitoring. These files are available to you to help in resolving problems encountered while using the products. IBM Software Support might request some or all of these files while investigating a problem you have reported.

For general problem determination information, see the *IBM Tivoli Monitoring Troubleshooting Guide*. For problem determination relating to the monitoring agent, see the *IBM Tivoli OMEGAMON XE on z/VM and Linux Troubleshooting Guide*.

Log files

The log files created by the IBM Tivoli Monitoring and the OMEGAMON XE monitoring agents contain messages and trace information about the events and processing being performed. IBM Tivoli Monitoring log files provide a complete record of system activity. These files are created when you start each of the IBM Tivoli Monitoring components. Table 7 lists the log files created by each component.

Table 7. Locations of various types of logs

Component	File name and path	Description
Tivoli Enterprise Monitoring Agent See "Definitions of variables for RAS1 logs" on page 74 for descriptions of the variables in the file names in column two.	<i>install_dir</i> /logs/ <i>hostname_PC_timestamp</i> .log	Traces activity on the monitoring agent. Note: Other logs, such as logs for collector processes and Take Action commands (if available), have a similar syntax and are located in this directory path.
Tivoli Enterprise Monitoring Server See "Definitions of variables for RAS1 logs" on page 74 for descriptions of the variables in the file names in column two.	On UNIX: <i>install_dir</i> /logs/ <i>hostname_ms_timestamp</i> .log On Windows: <i>install_dir</i> \logs\ <i>hostname_ms_HEXtimestamp- nn</i> .log	Traces activity on the monitoring server. Note: Trace logging is enabled by default. A configuration step is not required to enable this tracing.
Tivoli Enterprise Portal Server See "Definitions of variables for RAS1 logs" on page 74 for descriptions of the variables in the file names in column two.	On UNIX: <i>install_dir</i> /logs/ <i>hostname_cq_timestamp</i> .log On Windows: <i>install_dir</i> \logs	Traces activity on the portal server.
Tivoli Enterprise Portal browser client	On Windows: C:\Documents and Settings\Administrator\Application Data\Java\Deployment\log\plugin142.trace On UNIX: None.	
Tivoli Enterprise Portal desktop client	On UNIX: <i>install_dir</i> /logs/ <i>hostname_cj_timestamp</i> .log On Windows: <i>install_dir</i> \CNP\kcjerror.log <i>install_dir</i> \CNP\kcjras1.log	
IBM Tivoli Warehouse Proxy agent	On Windows: <i>install_dir</i> \logs\ <i>hostname_hd_timestamp</i> .log Not supported on UNIX computers.	

Table 7. Locations of various types of logs (continued)

Component	File name and path	Description
KVL CONFIG file is associated with the Take Action Command Processor REXX executable (KVL CMD EXEC) of the OMEGAMON XE on z/VM and Linux monitoring agent. You may optionally update the KVL CONFIG file to output the results of the Take Action commands to the KVL <i>nn</i> log file.	<p>KVL <i>nn</i></p> <p>where:</p> <p><i>nn</i> represents the incremental name for the log file, starting with 1. The number of log files generated depends on the LOG_COUNT parameter specified in the KVL CONFIG file. The log cycles through as many times as you specify in this setting. A maximum of five active log files can be retained. The default is 2.</p> <p>The log file is located on the A-disk of the Command Processor virtual machine.</p>	A log file is created the first time the Command Processor REXX executable file runs. The log file is updated as the CMS guest system processes the Take Action commands from the monitoring agent. The log file is a rolling log, containing information about all of the commands received by the CMS guest.

The messages issued by the OMEGAMON XE monitoring agent products are documented in the product documentation for the individual agents. The messages are listed in alphanumeric order by message number. The message number begins with a prefix that identifies the product or component. The messages associated with this monitoring agent are documented in the *IBM Tivoli OMEGAMON XE on z/VM and Linux Troubleshooting Guide*.

When you encounter a problem, first check the messages in the log files to determine if it is a problem in your environment or with a IBM Tivoli Monitoring product. If the problem seems to be a product defect, you might find a problem that matches your symptoms and a solution in the IBM Support Knowledge database. Refer to “Support for problem solving” on page 101 for details about using this support tool.

IBM Software Support might ask you to activate tracing so that the log files will collect additional information needed to resolve the problem. Some of the tracing options produce large amounts of trace information. Monitor the disk or spool space when activating tracing to prevent your disk or spool from reaching capacity. Return the trace settings to the default settings after the desired trace information has been collected.

Definitions of variables for RAS1 logs

The following definitions apply to the RAS1 logs:

- *hostname* is the name of the system hosting the product.
- *install_dir* represents the directory path where you installed the IBM Tivoli Monitoring component. *install_dir* can represent a path on the computer that hosts the monitoring server, the monitoring agent, or the portal server.
- *PC* is the two-character product code. The product code is v1 for the Tivoli OMEGAMON XE on z/VM and Linux monitoring agent. The product code is 1z for the IBM Tivoli Monitoring Agent for Linux OS agent.
- *HEXtimestamp* is a hexadecimal representation of the time at which the process was started.
- *nn* is the circular sequence in which logs are rotated. Ranges from 1-5, by default, though the first is always retained, since it includes configuration parameters.
- *timestamp* is a decimal representation of the time at which the process was started.

Log files for some of the components are explained in the sections that follow.

See the *IBM Tivoli Monitoring Troubleshooting Guide* for more information on the complete set of trace logs that are maintained on the monitoring server.

Note: When you communicate with IBM Software Support, you must capture and send the RAS1 log that matches any problem occurrence that you report.

Tivoli Enterprise Monitoring Server on Windows computers or on UNIX computers

The log files are created automatically when you start Tivoli Enterprise Monitoring Server on Windows computers or on UNIX computers. You can view the log file with any text editor.

When you investigate problems with Tivoli Enterprise Monitoring Server, use the Windows Event Viewer to check that the Tivoli Enterprise Monitoring Server started correctly and to look for errors.

Tivoli Enterprise Portal

The log files are created automatically when you start Tivoli Enterprise Portal. You can view the log files with any text editor.

In desktop mode, the log files are named kcjas1.log and kcjerror.log. kcjerror.log contains any errors that might have been written by the Sun Java libraries that are used by the Tivoli Enterprise Portal desktop client.

In browser mode, the log file is named plugin142.trace. You can change the level of tracing by using the **File > Trace Options...** window.

In addition, logon prompts and progress messages are displayed in the Logon window status bar. This area is also used to display error messages.

When you investigate problems with Tivoli Enterprise Portal, use the Windows Event Viewer to check that the Tivoli Enterprise Portal Server started correctly and to look for errors.

For more information about troubleshooting problems on Tivoli Enterprise Portal, refer to *IBM Tivoli Monitoring Troubleshooting Guide*.

Tivoli Enterprise Portal Server

The log files are created automatically when you start Tivoli Enterprise Portal Server. You can view the log files with any text editor.

You can change trace settings using the Manage Tivoli Enterprise Monitoring Services **Action > Advanced > Edit Trace Params...** window. You can also use the Service Console, accessible from the Tivoli Enterprise Portal Server using an Internet Explorer browser, to read logs and turn on traces for remote product diagnostics and configuration.

For more information about troubleshooting problems on Tivoli Enterprise Portal Server, refer to the *IBM Tivoli Monitoring Troubleshooting Guide*.

Tivoli Data Warehouse and the warehouse proxy

To view the Application Event Log for Tivoli Data Warehouse, start the Event Viewer by clicking **Start > Programs > Administrative Tools > Event Viewer**. Select **Application** from the **Log** pull-down menu.

In the warehouse proxy, you can set error tracing to capture additional error messages that can be helpful in detecting problems. Refer to the *IBM Tivoli Monitoring Troubleshooting Guide* for more information.

Part 4. Appendixes

Planning worksheets

Use the planning worksheets in this chapter to collect the information required. Complete the following worksheets before configuring the product.

The planning worksheets are as follows:

- “Worksheet: Your overall configuration”
- “Worksheet: Your monitoring agent configuration” on page 80
- “Worksheet: Planning communication protocols for the monitoring agent when the monitoring server is on a distributed system” on page 81
- “Worksheets: Information to gather when configuring your portal server on Windows or Linux” on page 82
- “Worksheets: Information to gather when configuring your monitoring server on a distributed system” on page 84
- “Worksheets: Information to gather when putting your hub monitoring server on a z/OS system” on page 86
- “Worksheets: Information to gather when configuring your portal desktop client on Windows or on Linux” on page 91
- “Specifying communication protocols between components” on page 93

Worksheet: Your overall configuration

Use the following worksheets to fill in your own overall OMEGAMON XE on z/VM and Linux configuration. Fill in the system name where you plan to install and configure each component, using “Understanding and designing your configuration” on page 16 as a guide.

Table 8. Worksheet for designing your overall configuration

IBM Tivoli Management Services component	Installation system name
Tivoli Enterprise Portal desktop client	Desktop client is located on (check one): <ul style="list-style-type: none"> • Windows • Intel Linux
Tivoli Enterprise Portal browser client	Windows
Tivoli Enterprise Portal Server	Tivoli Enterprise Portal Server is located on (check one): <ul style="list-style-type: none"> • Windows • Linux
Hub Tivoli Enterprise Monitoring Server	Hub Tivoli Enterprise Monitoring Server is located on (check one): <ul style="list-style-type: none"> • Windows server • Linux server • UNIX server • z/OS server <ul style="list-style-type: none"> – Hostname: _____ – IP address: _____
Remote Tivoli Enterprise Monitoring Servers	<ul style="list-style-type: none"> • Yes • No <p>If yes, indicate where you will put remote Tivoli Enterprise Monitoring Server or Servers:</p> <ul style="list-style-type: none"> • Windows server • Linux server • UNIX server • z/OS server <ul style="list-style-type: none"> – Hostname: _____ – IP address: _____

For complete information about operating system version support for each Tivoli Management Services component, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

Worksheet: Your monitoring agent configuration

Use the following worksheet to fill in your own OMEGAMON XE on z/VM and Linux monitoring agent configuration. Fill in the system name where you plan to install and configure each component.

Table 9. Worksheet for configuring your monitoring agent

Value	Description	Value for your configuration
LPAR name	The name assigned to the logical partition.	_____
Installation directory	Installation directory where you installed IBM Tivoli Monitoring and the monitoring agents.	
Encryption key	You are prompted for a 32-bit encryption key when you begin configuration of the portal server and other components on a distributed system. You can use the default key. Be sure to document the value you use for the key, because you must use the same key in configuring any hub monitoring server that communicates with this portal server.	
Agent to install	See the <i>IBM Tivoli Monitoring Installation and Setup Guide</i> for complete information about installing the monitoring agents.	Note: One agent instance is needed per monitored z/VM image. The name of the OMEGAMON XE on z/VM and Linux system: • Linux Guest user ID _____
Agent to install	Space is provided for the Tivoli Monitoring Agent for Linux OS system name in this worksheet. See the <i>IBM Tivoli Monitoring Installation and Setup Guide</i> for complete information about limitations and considerations in installing multiple monitoring agents. Note: The installation of this monitoring agent is only required if you intend to use dynamic workspace linking. See "Dynamic linking to cross-product workspaces" on page 9.	The name of the Tivoli Monitoring Agent for Linux OS system: • Linux Guest user ID _____
Agent product code or codes	Product code for each agent whose data you want to send to the monitoring server. • KVL for the OMEGAMON XE on z/VM and Linux monitoring agent • KLZ for the Tivoli Monitoring Agent for Linux OS (agent is only required for dynamic workspace linking)	• Product code _____ • Product code _____
Additional OMEGAMON XE on z/VM and Linux monitoring agent configuration requirements.	• To display the data collected by the Performance Toolkit, you need the name of the DCSS (the default name is PERFOUT). See "Step 3. Defining a DCSS on z/VM" on page 45. • To enable Take Action commands (optional), you need the name of the z/VM CMS guest system that is configured to receive commands from the Tivoli Enterprise Portal. Take Action commands execute CP, CMS, and REXX execs on the z/VM operating system. See "z/VM operating system software requirements for the OMEGAMON XE on z/VM and Linux monitoring agent" on page 22. See also "Step 11. Enabling Take Action commands (optional)" on page 53.	• Name of the DCSS defined in z/VM _____ • z/VM Conversational Monitor System (CMS) guest system where the Command Processor is running: • CMS guest system name _____
Monitoring server host name	The host name of the workstation where the hub monitoring server is installed. You'll need both the short host name (without the domain name), and the fully qualified host name of the monitoring server workstation (with the domain name).	
Communication protocols	See "Worksheet: Planning the communication protocols for the portal server" on page 83	
KDC_PARTITION	You must create or modify this file before implementing firewall support with the monitoring server and agents. This step can be performed at a later time.	• Use default key: _____ • Define your own key: _____
NIC interface name (Optional Primary Network Name)	To establish connectivity between the monitoring server and agents, you must specify an additional variable when configuring the monitoring server or agents.	
root user password		
User group name		
Optional user name		

For complete information about operating system version support for each Tivoli Management Services component, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

Worksheet: Planning communication protocols for the monitoring agent when the monitoring server is on a distributed system

Your monitoring agent must communicate with the monitoring server. Gather the following information to configure communication for your monitoring agent:

When you have configured your monitoring server on a distributed system, you must plan a communication protocol for the monitoring agent to send data from the monitoring agent to the hub monitoring server.

Make sure that at least one of the protocols you specify for the monitoring agent correspond to a protocol specified for the monitoring server. See “Worksheet: Planning communication protocols for the monitoring server on a distributed system” on page 85 for the values you specified for your monitoring server communication protocols.

You can select from the following protocols:

IP.PIPE

Uses the TCP/IP protocol for underlying communications. The IP.PIPE protocol is generally the best choice for Protocol 1 in firewall environments. It enables the Tivoli Enterprise Monitoring Server to communicate with the Tivoli Enterprise Portal Server on a Windows system, even if both the z/OS system and the Windows system are running behind firewalls.

IP.UDP

Uses the TCP/IP User Datagram Protocol (UDP).

IP.SPIPE

Secure IP.PIPE protocol.

Note: If you are running IBM Tivoli Monitoring, V6.1.x, you cannot specify the IPV6 protocol for the monitoring agent. If you are running IBM Tivoli Monitoring, V6.2.x, you can use the IPV6 protocol. See the technical note on the support page for this monitoring agent entitled *IBM Tivoli OMEGAMON XE on z/VM and Linux Post-Version 411 Support Update* for information on configuring the monitoring agent and the IBM Tivoli Monitoring shared technology components to use the IPV6 protocol. The support page is accessed at the following address:

<http://www-306.ibm.com/software/sysmgmt/products/support/IBMTivoliOMEGAMONonzVMLinux.html>

Use the following worksheet to gather information for specifying the communication protocols for the monitoring agent when the monitoring server is on a distributed system.

Table 10. Worksheet for specifying communication protocols for the monitoring agent when the monitoring server is on a distributed system

Value	Description	Value for your configuration
Monitoring server name (node ID)	Node ID of the hub monitoring server. Note that the Node ID is generally not the same as the host name. It is an arbitrary name assigned during Tivoli Enterprise Monitoring Server configuration. Find the Node ID as follows, depending on where the monitoring server is installed: <ul style="list-style-type: none"> On Windows systems, you can find the Node ID in Manage Tivoli Enterprise Monitoring Services. Right-click the Tivoli Enterprise Monitoring Server and select Browse Settings, and look for the value of TEMS_NODEID. On Linux and UNIX systems, you can find the value of TEMS_NODEID in the KBBENV file located in the \$CANDLEHOME/tables/cms_name subdirectory. On z/OS systems, you can find the value of TEMS_NODEID in this location: &rhilev.&sys.RKANPAR(KDSENV) 	
IP.PIPE or IP.SPIPE	For these protocols, you must gather the following values:	

Table 10. Worksheet for specifying communication protocols for the monitoring agent when the monitoring server is on a distributed system (continued)

Value	Description	Value for your configuration
Hostname	Host name of the system where the monitoring server is installed.	
Address	The IP address of the system where the monitoring server is installed.	
Port number	Same IP port number that you specified for the monitoring server. See "Worksheet: Planning communication protocols for the monitoring server on a distributed system" on page 85.	

Worksheets: Information to gather when configuring your portal server on Windows or Linux

If you are putting your portal server on Windows or Linux, fill out the tables below:

- "Worksheet: Information for configuring your portal server on Windows"
- "Worksheet: Information for configuring your portal server on Linux"
- "Worksheet: Planning the communication protocols for the portal server" on page 83

Worksheet: Information for configuring your portal server on Windows

Gather the following information for use in configuring your portal server on Windows:

Table 11. Worksheet for configuring your portal server on Windows

Value	Value for your configuration
Installation directory	
Encryption key used on the hub monitoring server	
Program folder	
Host name of the computer where you installed 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)	
Hub monitoring server host name	
Communication protocols	See "Worksheet: Planning the communication protocols for the portal server" on page 83

Worksheet: Information for configuring your portal server on Linux

Gather the following information for use in configuring your portal server on Linux:

Table 12. Worksheet for configuring your portal server on Linux

Value	Value for your configuration
Installation location	
Encryption key for the hub monitoring server	
Host name for the hub monitoring server	
NIC interface name (Primary Optional Network Name)	
DB2 instance name (default = db2inst1)	
DB2 administrator ID (default = db2inst1)	
DB2 administrator password	

Table 12. Worksheet for configuring your portal server on Linux (continued)

Value	Value for your configuration
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	
Communication protocols	See "Worksheet: Planning the communication protocols for the portal server"

Worksheet: Planning the communication protocols for the portal server

You must plan a communication protocol for the portal server to receive data from the monitoring server and communicate with portal servers. You can select one of the following protocol types:

IP.PIPE

Uses the TCP/IP protocol for underlying communications. The IP.PIPE protocol is generally the best choice for Protocol 1 in firewall environments. This protocol enables the Tivoli Enterprise Monitoring Server to communicate with the Tivoli Enterprise Portal Server on a Windows system, even if both the z/OS system and the Windows system are running behind firewalls.

IP.UDP

Uses the TCP/IP User Datagram Protocol (UDP).

IP.SPIPE

Secure IP.PIPE protocol.

SNA.PIPE

Uses the SNA Advanced Program-To-Program Communications (APPC). Because some zSeries OMEGAMON products require SNA, the SNA.PIPE protocol must be one of the protocols for Tivoli Enterprise Monitoring Server on z/OS. However, it need not be Protocol 1 (the highest-priority protocol).

Note: If you are running IBM Tivoli Monitoring, V6.1.x, you cannot specify the IPV6 protocol for the portal server. If you are running IBM Tivoli Monitoring, V6.2.x, you can use the IPV6 protocol. See the technical note on the support page for this monitoring agent entitled *IBM Tivoli OMEGAMON XE on z/VM and Linux Post-Version 411 Support Update* for information on configuring the monitoring agent and the IBM Tivoli Monitoring shared technology components to use the IPV6 protocol. The support page is accessed at the following address:

<http://www-306.ibm.com/software/sysmgmt/products/support/IBMTivoliOMEGAMONonzVMLinux.html>

Use the following worksheet to gather information for specifying the communication protocols for the portal server.

Table 13. Worksheet for specifying communication protocols for the portal server

Field	Description	Value for your configuration
IP.UDP Settings	Uses the TCP/IP User Datagram Protocol (UDP). Gather the values requested in the following rows:	
Hostname or IP Address	The host name or IP address of the monitoring server.	
Port # and/or Port Pools	The listening port for the hub monitoring server to use in communicating with the monitoring server, or the pool from which the port is to be selected. The default number is 1918.	
IP.PIPE Settings	Uses the TCP/IP protocol for underlying communications. The TCP/IP protocol is the best choice for Protocol 1 in a firewall environment. Gather the following values:	

Table 13. Worksheet for specifying communication protocols for the portal server (continued)

Field	Description	Value for your configuration
Hostname or IP Address	The host name or IP address of the monitoring server.	
Port Number	The listening port for the hub monitoring server to use in communicating with the monitoring agent. The default number is 1918.	
IP.SPIPE Settings	Secure IP.PIPE protocol. Gather the following values:	
Hostname or IP Address	The host name or IP address you plan to use for the monitoring agent.	
Port Number	The listening port for the hub monitoring server to use in communicating with the monitoring agent. The default port number is 3660.	
SNA Settings	Uses the SNA Advanced Program-To-Program Communications (APPC). Gather the values requested in the following rows:	
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. The default value is SNASOCKETS .	

Worksheets: Information to gather when configuring your monitoring server on a distributed system

If you are putting your monitoring server on a distributed system, fill out the tables below:

- “Worksheet: Information for configuring your hub monitoring server on a distributed system”
- “Worksheet: Planning communication protocols for the monitoring server on a distributed system” on page 85

If you are putting your monitoring server on z/OS, fill out the worksheets in “Worksheets: Information to gather when putting your hub monitoring server on a z/OS system” on page 86.

Note that all fields are required, unless otherwise indicated.

Worksheet: Information for configuring your hub monitoring server on a distributed system

If you are installing your hub monitoring server on a distributed system, you must gather the following information for use in the configuration process for each hub monitoring server.

Table 14. Worksheet for configuring your hub monitoring server on a distributed system

Value	Description	Value for your configuration
Host name of the computer	The host name of the workstation where the hub monitoring server is installed. You'll need both the short host name (without the domain name), and the fully qualified host name of the monitoring server workstation (with the domain name).	
IP Address	The IP address of the workstation where the hub monitoring server is installed.	
IBM Tivoli Monitoring installation directory	The directory in which you installed the IBM Tivoli Monitoring components.	
Encryption key	You'll be prompted for a 32-bit encryption key when you begin configuration of components on a distributed system. You can use the default key. Be sure to document the value you use for the key, because you must use the same key in configuring any monitoring server and the portal servers that communicate.	<ul style="list-style-type: none"> • Use default key: _____ • Define your own key: _____

Table 14. Worksheet for configuring your hub monitoring server on a distributed system (continued)

Value	Description	Value for your configuration
Monitoring server name / node ID	<p>Name (node ID) of the monitoring server. The default name for the hub monitoring server is HUB_host_name. For example, for host ITMSERV61, the default hub name is HUB_ITMSERV61.</p> <p>The Node ID is generally not the same as the host name. It is an arbitrary name assigned during monitoring server configuration.</p> <ul style="list-style-type: none"> On Windows systems, you can find the Node ID in Manage Tivoli Monitoring Services. Right-click the Tivoli Enterprise Monitoring Server and select Browse Settings, and look for the value of CMS_NODEID. On Linux and UNIX systems, you can find the value of CMS_NODEID in the KBBENV file located in the \$CANDLEHOME/tables/cms_name subdirectory. 	
NIC interface name ("Optional Primary Network Name")	To establish connectivity between the monitoring server and agents, you must specify an additional variable when configuring the monitoring server or agents.	
KDC_PARTITION	You must create or modify this file before implementing firewall support with the monitoring server and agents. This step can be performed at a later time.	
Agents for which to add application support data	Product code for each agent whose data you want to send to the monitoring server.	
Agents to add to the deployment depot	Monitoring agents you added to the deployment depot.	
Communication protocols	See "Worksheet: Planning communication protocols for the monitoring server on a distributed system"	

Worksheet: Planning communication protocols for the monitoring server on a distributed system

You must plan communication protocols for a monitoring server on a distributed system to send data to other components, such as remote monitoring servers and portal servers. *If you are configuring your monitoring server on a distributed system*, you can select from the following protocols:

IP.PIPE

Uses the TCP/IP protocol for underlying communications. The IP.PIPE protocol is generally the best choice for Protocol 1 in firewall environments.

IP.UDP

Also a TCP/IP protocol. Uses the User Datagram Protocol (UDP).

IP.SPIPE

Secure IP.PIPE protocol.

SNA.PIPE

Uses the SNA Advanced Program-To-Program Communications (APPC). Because some zSeries OMEGAMON products require SNA, the SNA.PIPE protocol must be one of the protocols for Tivoli Enterprise Monitoring Server on z/OS. However, it need not be Protocol 1 (the highest-priority protocol).

Note: If you are running IBM Tivoli Monitoring, V6.1.x, you cannot specify the IPV6 protocol for the monitoring server on a distributed system. If you are running IBM Tivoli Monitoring, V6.2.x, you can use the IPV6 protocol for the monitoring server on a distributed system. See the technical note on the support page for this monitoring agent entitled *IBM Tivoli OMEGAMON XE on z/VM and Linux Post-Version 411 Support Update* for information on configuring the monitoring agent and the IBM Tivoli Monitoring shared technology components to use the IPV6 protocol. The support page is accessed at the following address:

<http://www-306.ibm.com/software/sysmgmt/products/support/IBMTivoliOMEGAMONonzVMLinux.html>

Use the following worksheet to gather information for specifying the communication protocols for your hub and remote monitoring servers on distributed systems.

Table 15. Worksheet for specifying communication protocols for a monitoring server on a distributed system

Field	Description	Value for your configuration
IP.UDP Settings	Uses the TCP/IP User Datagram Protocol (UDP). Gather the values required in the following rows:	
Hostname or IP Address	The host name or IP address of the system where you plan to install the monitoring agent.	
Port # and/or Port Pools	The listening port for the hub monitoring server to use in communicating with the monitoring agent, or the pool from which the port is to be selected. The default number is 1918.	
IP.PIPE Settings	Uses the TCP/IP protocol for underlying communications. This is the best choice for Protocol 1 in a firewall environment. Gather the following values:	
Hostname or IP Address	The host name or IP address you plan to use for the monitoring agent.	
Port Number	The listening port for the hub monitoring server to use in communicating with the monitoring agent. The default port number is 1918.	
IP.SPIPE Settings	Secure IP.PIPE protocol. Gather the following values:	
Hostname or IP Address	The host name or IP address you plan to use for the monitoring agent.	
Port Number	The listening port for the hub monitoring server to use in communicating with the monitoring agent. The default port number is 3660.	
SNA Settings	Uses the SNA Advanced Program-To-Program Communications (APPC). Gather the following values:	
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. The default value is SNASOCKETS .	

Worksheets: Information to gather when putting your hub monitoring server on a z/OS system

If you are putting your hub monitoring server on a z/OS system, fill out the tables below:

- “Worksheet for configuring your hub monitoring server on a z/OS system”
- “Worksheet for configuring your communications protocols for a hub monitoring server on a z/OS system” on page 88

For detailed information on configuring your monitoring server on z/OS, see *Configuring IBM Tivoli Enterprise Monitoring Server on z/OS*.

Worksheet for configuring your hub monitoring server on a z/OS system

Note that all fields are required, unless otherwise indicated.

Table 16. Worksheet for configuring the hub monitoring server on z/OS

Value	Description	Value for your configuration
Fully qualified host name of the z/OS system where the hub monitoring server is installed	To obtain the host name, enter TSO HOMETEST at the command line of the z/OS system to be monitored and use the first qualifier of the TCP hostname. For example, sys is the first qualifier of a fully qualified TCP hostname sys.ibm.com.	
HTTP server port number	Accept the default value of 1920. This field is required for the SOAP Server, which must be enabled for a hub monitoring server on z/OS.	
LU6.2 logmode settings: You must associate a SNA logmode with a monitoring server on z/OS. You can either use an existing logmode or create a new one. Gather the following logmode information for the configuration process of each monitoring server on z/OS:		

Table 16. Worksheet for configuring the hub monitoring server on z/OS (continued)

Value	Description	Value for your configuration
LU6.2 logmode name	Name of the LU6.2 logmode defined for use by the monitoring server. The default value is CANCTDCS .	
Logmode table name	Name of the logmode table that contains the LU6.2 logmode. The default name is KDSMTAB1 .	
VTAMLIB load library	Name of the system library used to contain VTAM® logmode tables. This is usually SYS1.VTAMLIB. You can specify any load library if you do not want to update you VTAMLIB directly.	
VTAM® macro library	Name of the system library that contains the VTAM macros. This is usually SYS1.SISTMAC.	
Configuration value settings: Gather the following information about your monitoring server on z/OS:		
Tivoli Enterprise Monitoring Server started task name	Name of the started task (procedure name) for the monitoring server. This value should be eight characters, maximum.	
Hub or remote Tivoli Enterprise Monitoring Server	Indicate whether this is a hub or remote monitoring server.	
Is z/OS Integrated Cryptographic Service Facility (ICSF) installed?	You must have ICSF installed on the z/OS system where you install your hub monitoring server. Without ICSF installed, the monitoring server cannot connect to the Tivoli Enterprise Portal Server. If you want to install your hub monitoring server on a z/OS system that does not have ICSF installed, you must use a workaround to ensure communication between the monitoring server on z/OS and the Tivoli Enterprise Portal Server..	<ul style="list-style-type: none"> • Yes • No
ICSF load library	If ICSF is installed and configured on the z/OS system, specify the load library that contains the CSNB* modules used for password encryption.	
Encryption key	You are prompted for a 32-bit ICSF encryption key. You can use the default key. Be sure to document the value you use for the key, because you must use the same key during the installation of any components that communicate with this monitoring server.	<ul style="list-style-type: none"> • Use default key: _____ • Define your own key: _____
Enable Web Services SOAP Server	The Web Services SOAP Server must be enabled for a hub monitoring server. You must accept the default value of Y for the Enable Web Services SOAP Server field if you are configuring a hub monitoring server.	
Language locale	Specify a numeric value (1-36) representing the language and region. For example, specify 1 for United States English. For a list of the language locale values, press F1 in the Configuration Tool panel where the prompt is displayed.	
VTAM network ID	A VTAM network ID is required for any monitoring server on z/OS.	

Worksheet for configuring your communications protocols for a hub monitoring server on a z/OS system

Table 17. Worksheet for configuring the communications protocols for a hub monitoring server on z/OS

Value	Description	Value for your configuration
<p>Communications protocols for the monitoring server on z/OS:</p> <p>You will specify the communications protocols for the monitoring server in one of the following:</p> <ul style="list-style-type: none"> When you put the monitoring server and monitoring agent in the same address space in z/OS When you put the monitoring server and monitoring agent in different address spaces on z/OS <p>See the configuration guide for the z/OS monitoring agent for details.</p>	<p>You can choose from all the IP protocols we show in the list below. For the monitoring server and monitoring agent to communicate, at least one of the monitoring server protocols must be the same as one of the monitoring agent protocols. In addition, you must specify SNA.PIPE as one of the protocols for a Tivoli Enterprise Monitoring Server on z/OS. However, it need not be Protocol 1 (the highest-priority protocol).</p> <p>For a hub monitoring server on z/OS, you must specify a TCP/IP protocol (IP.PIPE or IP.UDP) as one of your protocols, for use by the Web Services SOAP Server, which must be enabled.</p> <p>Choose from the following protocols:</p> <p>IP.PIPE Uses the TCP/IP protocol for underlying communications. The IP.PIPE protocol is generally the best choice for Protocol 1 in firewall environments. It enables the Tivoli Enterprise Monitoring Server to communicate with the Tivoli Enterprise Portal Server on a Windows system, even if both the z/OS system and the Windows system are running behind firewalls.</p> <p>IP.UDP Also a TCP/IP protocol. Uses the User Datagram Protocol (UDP).</p> <p>IP6.PIPE IP.PIPE protocol with IPV6 installed and operational. This protocol is available only for a monitoring server on a z/OS system at release level V1R7 or higher.</p> <p>IP6.UDP IP.UDP protocol with IPV6 installed and operational. This protocol is available only for a monitoring server on a z/OS system at release level V1R7 or higher.</p> <p>IP.SPIPE Secure IP.PIPE protocol. This protocol is available only for a monitoring server on a z/OS system at release level V1R7 or higher.</p> <p>IP6.SPIPE Secure IP.PIPE for IPV6. This protocol is available only for a monitoring server on a z/OS system at release level V1R7 or higher with IPV6 installed and operational.</p> <p>SNA.PIPE Uses the SNA Advanced Program-To-Program Communications (APPC).</p> <p>SNA.PIPE Uses the SNA Advanced Program-To-Program Communications (APPC).</p>	<ul style="list-style-type: none"> Protocol 1 _____ Protocol 2 _____ Protocol 3 _____
<p>Settings for IP. and IP6. If you select any of the IP. or IP6. protocols, you must gather the following values:</p>		
Hostname	<p>The TCP/IP host name of the z/OS system to which the Tivoli Enterprise Monitoring Server connects to receive data from the monitoring agent.</p> <p>To obtain the host name, enter TS0 HOMETEST at the command line of the z/OS system to be monitored and use the first qualifier of the TCP hostname. For example, sys is the first qualifier of a fully qualified TCP hostname sys.ibm.com.</p>	
Address	<p>The IP address of the z/OS system to which the Tivoli Enterprise Monitoring Server connects to receive data from the monitoring agent.</p> <p>To obtain the IP Address, enter TS0 HOMETEST at the command line of the z/OS system to be monitored.</p>	
Started task	<p>The started task name of the TCP/IP server. You can specify * to allow the IP stack to dynamically find the TCP/IP image. * is the suggested value for the started task.</p>	

Table 17. Worksheet for configuring the communications protocols for a hub monitoring server on z/OS (continued)

Value	Description	Value for your configuration
HTTP server port number	Accept the default value of 1920. This field is required for the SOAP Server, which must be enabled for a hub monitoring server on z/OS.	
Access TEMS list via SOAP Server?	Accept the default value of Y. The Web Services SOAP Server must be enabled for a hub monitoring server on z/OS.	
Address translation	Specify Y to configure IP.PIPE support for communication across firewalls using address translation.	
Partition name	Specify the partition name that identifies the monitoring server (namespace) relative to the firewall(s) used for address translation.	
SNA.PIPE settings	If you pick the SNA.PIPE protocol, you must gather the following values:	
Applid prefix	Specify the applid prefix you want for all the VTAM applids required by the monitoring server. These applids begin with a prefix, and end with a unique applid value. The applids are contained in the VTAM major node. The default is CTDDS.	
Network ID	Specify the identifier for your VTAM network. You can locate this value on the NETID parameter in the VTAMLST startup member ATCSTRnn.	
<p>Communications protocols for the monitoring agent</p> <p>You will specify the communications protocol for the monitoring agent in a different address space than the monitoring server. See the configuration guide for the z/OS monitoring agent for details.</p>	<ul style="list-style-type: none"> • If you have configured the hub monitoring server and monitoring agent in the same address space, you can skip this section - you do not need to plan communication protocols for the monitoring agent. • If you configure the monitoring server and monitoring agent in different address space, you must plan communication protocols for the monitoring agent to send data to the hub monitoring server. Be sure to specify the same protocols you specified for the monitoring server. <p>If you specify a communications protocol for a monitoring agent in a different address space than the monitoring service, make sure that at least one of the protocols you specify for the monitoring agent is the same as one of the protocols specified for the monitoring server.</p> <p>You can select one of the following for each protocol:</p> <p>IP.PIPE Uses the TCP/IP protocol for underlying communications. The IP.PIPE protocol is generally the best choice for Protocol 1 in firewall environments. It enables the Tivoli Enterprise Monitoring Server to communicate with the Tivoli Enterprise Portal Server on a Windows system, even if both the z/OS system and the Windows system are running behind firewalls.</p> <p>For an IP.PIPE protocol, specify IPPPIPE in the Configuration Tool.</p> <p>IP.UDP Uses the TCP/IP User Datagram Protocol (UDP).</p> <p>For an IP.UDP protocol, specify IP in the Configuration Tool.</p> <p>SNA.PIPE Uses the SNA Advanced Program-To-Program Communications (APPC). Because some zSeries OMEGAMON products require SNA, it must be one of the protocols for Tivoli Enterprise Monitoring Server on z/OS. However, it need not be Protocol 1 (the highest-priority protocol).</p> <p>For a SNA.PIPE protocol, specify SNA in the Configuration Tool.</p>	<ul style="list-style-type: none"> • Protocol 1 _____ • Protocol 2 _____ • Protocol 3 _____

Table 17. Worksheet for configuring the communications protocols for a hub monitoring server on z/OS (continued)

Value	Description	Value for your configuration
TEMS name (node ID)	<p>Node ID of the hub monitoring server. Note that the node ID is generally not the same as the host name. It is an arbitrary name assigned during Tivoli Enterprise Monitoring Server configuration. Find the Node ID as follows, depending on where the monitoring server is installed:</p> <ul style="list-style-type: none"> On Windows systems, you can find the Node ID in Manage Tivoli Monitoring Services. Right-click the Tivoli Enterprise Monitoring Server and select Browse Settings, and look for the value of CMS_NODEID. On Linux and UNIX systems, you can find the value of CMS_NODEID in the KBBENV file located in the \$CANDLEHOME/tables/cms_name subdirectory. On z/OS systems, you can find the value of CMS_NODEID in this location: &rhilev.&sys.RKANPAR(KDSENV) 	
IPPIPE or IP: If you pick IPPIPE or IP, you must gather the following values:		
Hostname	The fully qualified TCP/IP host name of the z/OS system where the monitoring agent resides. To obtain the host name and IP address values, enter TSO HOMETEST at the command line on the z/OS system where the monitoring agent is installed.	
Address	The IP address of the z/OS system where the monitoring server is installed.	
Started task	The started task name of the TCP/IP server. You can specify * to allow the IP stack to dynamically find the TCP/IP image. * is the suggested value for the started task.	
Address translation	For IP.PIPE, specify Y to configure support for communication across firewalls using address translation.	
Partition name	For IP.PIPE, specify the partition name that identifies the monitoring server (namespace) relative to the firewalls used for address translation. Note that the monitoring server that the agent connects to must have a corresponding partition reference entry.	
SNA settings		
If you select the SNA protocol, you must gather the following value:		
Applid prefix	Specify the applid prefix you want for all the VTAM applids required by the monitoring server.	
<p>Communication protocol for the portal server</p> <p>You will specify the communications protocols for the portal server in one of the following:</p> <ul style="list-style-type: none"> When you put the monitoring server and z/OS monitoring agent in the same address space in z/OS When you put the monitoring server and z/OS monitoring agent in different address spaces on z/OS 	<p>You must plan a communication protocol for the portal server to receive data from the monitoring server and communicate with remote portal servers. You can select one of the following for each protocol:</p> <p>IP.PIPE Uses the TCP/IP protocol for underlying communications. The IP.PIPE protocol is generally the best choice for Protocol 1 in firewall environments. It enables the Tivoli Enterprise Monitoring Server to communicate with the Tivoli Enterprise Portal Server on a Windows system, even if both the z/OS system and the Windows system are running behind firewalls.</p> <p>IP.UDP Uses the TCP/IP User Datagram Protocol (UDP).</p> <p>SNA.PIPE Uses the SNA Advanced Program-To-Program Communications (APPC). Because some zSeries OMEGAMON products require SNA, it must be one of the protocols for Tivoli Enterprise Monitoring Server on z/OS. However, it need not be Protocol 1 (the highest-priority protocol).</p> <p>Use the following worksheet to gather information for specifying the communication protocols for the portal server:</p>	<ul style="list-style-type: none"> Protocol 1 _____ Protocol 2 _____ Protocol 3 _____
IP.UDP Settings: Uses the TCP/IP User Datagram Protocol (UDP). If you pick IP.UDP, gather the following values:		
Hostname or IP Address	The host name or IP address of the monitoring server.	
Port # and/or Port Pools	The listening port for the hub monitoring server to use in communicating with the monitoring server, or the pool from which the port is to be selected. The default number is 1918.	

Table 17. Worksheet for configuring the communications protocols for a hub monitoring server on z/OS (continued)

Value	Description	Value for your configuration
IP.PIPE Settings: Uses the TCP/IP protocol for underlying communications. This is the best choice for Protocol 1 in a firewall environment. If you pick IP.PIPE, gather the following values:		
Hostname or IP Address	The host name or IP address of the monitoring server.	
Port Number	The listening port for the hub monitoring server to use in communicating with the monitoring agent. The default number is 1918.	
SNA Settings: Uses the SNA Advanced Program-To-Program Communications (APPC). If you pick SNA (VTAM), gather the following values:		
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. The default value is SNASOCKETS .	

Worksheets: Information to gather when configuring your portal desktop client on Windows or on Linux

If you are putting your portal desktop client on either Windows or Linux, fill out the tables below:

- “Worksheet: Information for configuring your portal desktop client on Windows”
- “Worksheet: Information for configuring your portal desktop client on Linux”
- “Worksheet: Planning the communication protocols for the portal desktop client on a Windows system” on page 92

Worksheet: Information for configuring your portal desktop client on Windows

Gather the following information for use in configuring your portal desktop client on Windows:

Table 18. Worksheet for configuring your portal desktop client on Windows

Value	Value for your configuration
Installation directory	
Encryption key for the monitoring server	
Program folder name	
Host name of the computer where you installed the portal server	
Hub monitoring server host name	
Communication protocols	See “Worksheet: Planning the communication protocols for the portal desktop client on a Windows system” on page 92

Worksheet: Information for configuring your portal desktop client on Linux

Gather the following information for use in configuring your portal desktop client on Linux:

Table 19. Worksheet for configuring your portal desktop client on Linux

Value	Value for your configuration
Installation directory	
Encryption key for the monitoring server	
Portal server host name	
TCP/IP network services host name	

Worksheet: Planning the communication protocols for the portal desktop client on a Windows system

You must plan a communication protocol for the portal desktop client to receive data from the monitoring server. You can select one of the following protocol types:

IP.PIPE

Uses the TCP/IP protocol for underlying communications. The IP.PIPE protocol is generally the best choice for Protocol 1 in firewall environments. This protocol enables the Tivoli Enterprise Monitoring Server to communicate with the Tivoli Enterprise Portal desktop client running on Windows.

IP.UDP

Uses the TCP/IP User Datagram Protocol (UDP).

IP.SPIPE

Secure IP.PIPE protocol.

SNA.PIPE

Uses the SNA Advanced Program-To-Program Communications (APPC). Because some zSeries OMEGAMON products require SNA, the SNA.PIPE protocol must be one of the protocols for Tivoli Enterprise Monitoring Server on z/OS. However, it need not be Protocol 1 (the highest-priority protocol).

Note: If you are running IBM Tivoli Monitoring, V6.1.x, you cannot specify the IPV6 protocol for the portal desktop client on Windows. If you are running IBM Tivoli Monitoring, V6.2.x, you can use the IPV6 protocol. See the technical note on the support page for this monitoring agent entitled *IBM Tivoli OMEGAMON XE on z/VM and Linux Post-Version 411 Support Update* for information on configuring the monitoring agent and the IBM Tivoli Monitoring shared technology components to use the IPV6 protocol. The support page is accessed at the following address:

<http://www-306.ibm.com/software/sysmgmt/products/support/IBMTivoliOMEGAMONonzVMLinux.html>

Use the following worksheet to gather information for specifying the communication protocols for the portal desktop client on a Windows system.

Table 20. Worksheet for specifying communication protocols for the portal desktop client

Field	Description	Value for your configuration
IP.UDP Settings	Uses the TCP/IP User Datagram Protocol (UDP). Gather the values requested in the following rows:	
Hostname or IP Address	The host name or IP address of the monitoring server.	
Port # and/or Port Pools	The listening port for the hub monitoring server to use in communicating with the monitoring server, or the pool from which the port is to be selected. The default number is 1918.	
IP.PIPE Settings	Uses the TCP/IP protocol for underlying communications. The TCP/IP protocol is the best choice for Protocol 1 in a firewall environment. Gather the following values:	
Hostname or IP Address	The host name or IP address of the monitoring server.	
Port Number	The listening port for the hub monitoring server to use in communicating with the monitoring agent. The default number is 1918.	
IP.SPIPE Settings	Secure IP.PIPE protocol. Gather the following values:	
Hostname or IP Address	The host name or IP address you plan to use for the monitoring agent.	
Port Number	The listening port for the hub monitoring server to use in communicating with the monitoring agent. The default port number is 3660.	
SNA Settings	Uses the SNA Advanced Program-To-Program Communications (APPC). Gather the values requested in the following rows:	
Network Name	The SNA network identifier for your location.	

Table 20. Worksheet for specifying communication protocols for the portal desktop client (continued)

Field	Description	Value for your configuration
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. The default value is SNASOCKETS .	

Specifying communication protocols between components

IBM Tivoli Monitoring is built on communication between the components: The monitoring agent sends data to the monitoring server, which in turn routes it to the portal server, which sends it to the portal client for users to look at. As you can see in “Understanding and designing your configuration” on page 16, the arrows between the components represent communication protocols. You need to plan the following communication protocols:

Table 21. Summary: Planning communication protocols for IBM Tivoli Monitoring components

Where protocol is specified	Communications	Worksheet or document
Hub monitoring server on z/OS	Communicating with the monitoring agents, remote monitoring servers, and portal servers	“Worksheet for configuring your hub monitoring server on a z/OS system” on page 86. For detailed information on configuring the monitoring server on z/OS, see <i>Configuring IBM Tivoli Enterprise Monitoring Server on z/OS</i> .
Remote monitoring servers on z/OS	Communicating with the monitoring agents and with the hub monitoring server	
Hub monitoring server on a distributed system	Communicating with the monitoring agents, remote monitoring servers, and portal servers	“Worksheet: Planning communication protocols for the monitoring server on a distributed system” on page 85
Remote monitoring server on a distributed system	Communicating with the monitoring agents and with the hub monitoring server	
Monitoring agent when the monitoring server is configured on a distributed system	Communicating with the monitoring server	“Worksheet: Planning communication protocols for the monitoring agent when the monitoring server is on a distributed system” on page 81
Portal server	Communicating with the hub monitoring server	“Worksheet: Planning the communication protocols for the portal server” on page 83

You can select more than one protocol for each connection between components. The system then tries the protocols in the order you specify.

IP.PIPE

Uses the TCP/IP protocol for underlying communications. The IP.PIPE protocol is generally the best choice for Protocol 1 in firewall environments. It enables the Tivoli Enterprise Monitoring Server to communicate with the Tivoli Enterprise Portal Server on a Windows system, even if both the z/OS system and the Windows system are running behind firewalls.

IP.UDP

Uses the User Datagram Protocol (UDP) protocol.

IP6.PIPE

IP.PIPE protocol with IPV6 installed and operational.

IP6.UDP

IP.UDP protocol with IPV6 installed and operational.

IP.SPIPE

Secure IP.PIPE protocol. You can specify the IP.SPIPE if the minimum z/OS version on this mainframe system is V1.7.

IP6.SPIPE

Secure IP.PIPE for IPV6. You can specify IP6.SPIPE if the minimum z/OS version on this mainframe system is V1.7 and IPV6 is installed and operational.

SNA.PIPE

Uses the SNA Advanced Program-To-Program Communications (APPC). Because some zSeries

OMEGAMON products require SNA, it must be one of the protocols for Tivoli Enterprise Monitoring Server on z/OS. However, it need not be Protocol 1 (the highest-priority protocol).

Finding the information you need

This section provides a description of the kinds of tasks you perform as a user of the IBM Tivoli Monitoring and the OMEGAMON XE on z/VM and Linux agent, with pointers to the locations of the information required to perform these tasks. This information map applies to the following levels of documentation:

- OMEGAMON XE on z/VM and Linux information, V4.1.0 and above
- IBM Tivoli Monitoring: V6.1.0 for the Tivoli Enterprise Monitoring Server and V6.1.0 for Tivoli Enterprise Portal and Tivoli Enterprise Portal Server and above

Planning tasks

Table 22 lists the location of information required to plan the IBM Tivoli Monitoring and the OMEGAMON XE on z/VM and Linux monitoring agent.

Table 22. Planning tasks for IBM Tivoli Monitoring and monitoring agents

Component	Information contents	Information location
Planning for the IBM Tivoli Monitoring on Windows or UNIX	<ul style="list-style-type: none"> • Overall IBM Tivoli Monitoring information • Detailed information for installing components on Windows or UNIX systems • Installation planning (including firewall planning) • Prerequisites (hardware and software) and requirements 	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>
Planning for the Tivoli Enterprise Monitoring Server on z/OS	<ul style="list-style-type: none"> • General planning tasks • Security 	<i>Configuring IBM Tivoli Enterprise Monitoring Server on z/OS</i>
Planning for Tivoli Enterprise Portal and Tivoli Enterprise Portal Server	<ul style="list-style-type: none"> • Upgrading and migrating from previous releases • Information you need to have ready before you start the InstallShield wizard • Host names for TCP/IP network services • Order of installation and upgrades • Windows tasks <ul style="list-style-type: none"> – Windows logon ID – Host file – Downloading a JRE – Installing DB2 software • Guidelines for UNIX installations • UNIX tasks <ul style="list-style-type: none"> – SNA communications protocols – Creating a Tivoli account for installing and maintaining \$candlehome – Transferring files to UNIX using FTP 	<i>IBM Tivoli Monitoring Installation and Setup Guide and IBM Tivoli OMEGAMON XE and Tivoli Management Services on z/OS Upgrade Guide</i>
Planning for Tivoli Data Warehouse and historical data collection	<ul style="list-style-type: none"> • Developing a strategy for historical data collection • Collection rules • Warehousing data • Conversion programs • Columns and descriptions added automatically 	<i>IBM Tivoli Monitoring Installation and Setup Guide and IBM Tivoli Monitoring Administrator's Guide</i>
Planning for OMEGAMON XE on z/VM and Linux	<ul style="list-style-type: none"> • Understanding of product function and components • Knowledge of prerequisites for the platform and product • Installation flow 	Chapter 2, "Planning your OMEGAMON XE on z/VM and Linux configuration," on page 15

Understanding temporary defects, limitations, and workarounds

To understand all the information that might have changed since your installation media was created, look at the documents and Web sites shown in Table 23.

Table 23. Changes to product levels since the installation media was created for the IBM Tivoli Monitoring and monitoring agents

Component	Information contents	Information location
Preventative Service Planning (PSP) information is created by IBM software support	Information in the PSP bucket is identified by upgrade number and subset number. The subset number is the same as the FMID designation that SMP/E uses to install the product. For more information, refer to the program directory for the component or monitoring agent you are installing.	PSP bucket for IBM Tivoli Monitoring: http://techsupport.services.ibm.com
Late-breaking information	Defects, limitations, and workarounds for the product and the platform	<i>IBM Tivoli OMEGAMON XE on z/VM and Linux Readme</i>
Fixes for Windows-based components or UNIX-based components	Fix packs and service packs for workstation-based components	<ol style="list-style-type: none"> 1. From www.ibm.com, select Support and Downloads. 2. On the Support and Downloads page, find the block labeled Support by Product and click on Software. 3. From the Software Support page, select your product by name using the A to Z index and the resulting pop-up window. Remember to include "IBM Tivoli" in the name. 4. From the software support page for your product, scroll down to the Download heading and click on the appropriate topic, such as Latest Fixpack. Download and install the fix pack using the information on the page.

Upgrading from an earlier release

If you had an earlier version of the IBM Tivoli Monitoring or of the IBM Tivoli Monitoring Agent for Linux OS monitoring agent installed, refer to Table 24 for instructions on how to migrate from earlier releases. This is the first release of OMEGAMON XE on z/VM and Linux.

Table 24. Upgrade tasks for IBM Tivoli Monitoring and the monitoring agents

Component	Information contents	Information location
Upgrade any platform components installed on Windows or UNIX systems	<ul style="list-style-type: none"> • Finding out what fix packs are installed • Preserving a previous version on Windows • Upgrading a remote Tivoli Enterprise Monitoring Server while retaining a prior version on the hub Tivoli Enterprise Monitoring Server • Affinity considerations on Windows • UNIX directory structure changes • Stopping components • Database requirements • Adding application support • Defining a new \$candlehome location 	<i>IBM Tivoli Monitoring Installation and Setup Guide</i> and <i>IBM Tivoli OMEGAMON XE and Tivoli Management Services on z/OS Upgrade Guide</i> and Fix Pack Readme documentation
Migrating your database from MSDE or Microsoft SQL Server to IBM DB2 Universal Database	How the database migration takes place	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>
Replicating your Tivoli Enterprise Portal environment	<ul style="list-style-type: none"> • Replicating your Tivoli Enterprise Portal setup and customizations, which is the same procedure as moving from a test environment to a production environment and for backing up 	<i>IBM Tivoli Monitoring Administrator's Guide</i>

Installing and uninstalling

Table 25 lists the location of information required to install the components of the IBM Tivoli Monitoring, the OMEGAMON XE on z/VM and Linux monitoring agent, and the Tivoli Monitoring Agent for Linux OS, if you have installed this monitoring agent.

Table 25. Installing tasks for the IBM Tivoli Monitoring and monitoring agents

Component	Information contents	Information location
Installing Tivoli Enterprise Portal, the Tivoli Enterprise Portal Server, the Tivoli Data Warehouse and Tivoli Enterprise Monitoring Server on Windows or UNIX	Topics such as: <ul style="list-style-type: none"> • Starting the InstallShield wizard • Selecting products and copying files • Installing the hub Tivoli Enterprise Monitoring Server • Installing the remote Tivoli Enterprise Monitoring Server • Installing the Tivoli Enterprise Portal Server • Installing and configuring the monitoring agents • Installing the Tivoli Enterprise Portal desktop client • Installing support for the monitoring agents on the Tivoli Enterprise Monitoring Server • Installing the language packs • Starting and stopping the Tivoli Enterprise Portal client 	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>
Uninstalling distributed components	General uninstall instructions for Windows, Linux, and UNIX operating systems	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>

Configuring

Table 26 lists the location of information required to configure the components of the IBM Tivoli Monitoring and the OMEGAMON XE on z/VM and Linux monitoring agent.

Table 26. Configuration tasks for the IBM Tivoli Monitoring and monitoring agents

Component	Information contents	Information location
Configuring Tivoli Enterprise Portal	Configuration tasks associated with Tivoli Enterprise Portal	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>
Configuring the IBM Tivoli Monitoring components on Windows or UNIX	<ul style="list-style-type: none"> • Starting Manage Tivoli Enterprise Monitoring Services • Changing the configuration of the Tivoli Enterprise Monitoring Server • Configuring monitoring agents and the Tivoli Enterprise Monitoring Server • Starting or stopping components • Configuring user security • Configuring failover support • Adding application support to the Tivoli Enterprise Monitoring Server • Configuring the heartbeat interval 	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>
Configuring the warehouse proxy	<ul style="list-style-type: none"> • Special requirements • Setting up the ODBC connections • Configuring and registering the warehouse proxy • Configuring the ODBC data source • Error reporting • Behaviors 	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>
Configuring historical data collection and the Tivoli Data Warehouse	<ul style="list-style-type: none"> • Prerequisites • Naming the history tables • Attribute formatting • Logging 	<i>IBM Tivoli Monitoring Administrator's Guide</i>
Configuring Summarization and Pruning of historical data	<ul style="list-style-type: none"> • Planning to summarize and prune your collected data • Capacity planning suggestions for historical data collection on your Tivoli data warehouse 	<i>IBM Tivoli Monitoring Administrator's Guide</i>

Table 26. Configuration tasks for the IBM Tivoli Monitoring and monitoring agents (continued)

Component	Information contents	Information location
Configuring OMEGAMON XE on z/VM and Linux	<ul style="list-style-type: none"> Enabling the collection of data Estimating the size of the PERFOUT DCSS Defining a DCSS on z/VM Configuring the DCSS device driver on the Linux guest Loading the DCSS device driver Adding the PERFOUT DCSS to your Linux guest Loading the PERFOUT DCSS at startup time Enabling Take Action commands (optional) Defining user IDs 	Chapter 5, "Configuration required for the OMEGAMON XE on z/VM and Linux monitoring agent," on page 37 and Chapter 6, "Defining user IDs and security," on page 61
Configuring Tivoli Monitoring Agent for Linux OS to enable dynamic workspace linking	Configuration tasks required to enable dynamic workspace linking	"Dynamic linking to cross-product workspaces" on page 9 and "Step 10. Enabling dynamic workspace linking" on page 52

Tuning

Table 26 on page 97 lists the location of information required to tune the components of the IBM Tivoli Monitoring and the OMEGAMON XE on z/VM and Linux monitoring agent.

Table 27. Tuning the IBM Tivoli Monitoring and monitoring agents

Component	Information contents	Information location
Tuning Tivoli Enterprise Portal	Tivoli Enterprise Portal	<i>IBM Tivoli Monitoring Enterprise Portal User's Guide</i>
Tuning IBM Tivoli Monitoring	IBM Tivoli Monitoring	<i>IBM Tivoli Monitoring Installation and Setup Guide</i>
Tuning OMEGAMON XE on z/VM and Linux	OMEGAMON XE on z/VM and Linux	<i>IBM Tivoli OMEGAMON XE on z/VM and Linux Troubleshooting Guide</i> and Chapter 7, "Performance and storage considerations," on page 65

Administration

Administrative tasks for the IBM Tivoli Monitoring and the monitoring agents are found in Table 28.

Table 28. Administrative tasks for the IBM Tivoli Monitoring and monitoring agents

Component	Information contents	Information location
Administrative tasks associated with Tivoli Enterprise Portal and Tivoli Enterprise Portal Server	<ul style="list-style-type: none"> Setting up the environment, such as defining global parameters and variables and adding platforms to the Navigator Adding, editing, and removing users and assigning user permissions Setting up and using OMEGAMON Web services 	<i>IBM Tivoli Monitoring Administrator's Guide</i>
Administrative tasks associated with Tivoli Enterprise Monitoring Server on Windows or UNIX	Administrative tasks	<i>IBM Tivoli Monitoring Administrator's Guide</i>
Administrative tasks associated with OMEGAMON XE on z/VM and Linux	Adding z/VM user IDs	Chapter 6, "Defining user IDs and security," on page 61

Diagnosis

Table 29 lists the location of information required to install the components of the IBM Tivoli Monitoring and the OMEGAMON XE on z/VM and Linux monitoring agent.

Table 29. Diagnosis tasks for the IBM Tivoli Monitoring and monitoring agents

Component	Information contents	Information location
Diagnosing problems with OMEGAMON XE on z/VM and Linux	Messages for the OMEGAMON XE on z/VM and Linux monitoring agent <ul style="list-style-type: none"> • Where to look for product logs • Message prefixes that indicate messages from OMEGAMON XE on z/VM and Linux • Message explanations and action 	<i>IBM Tivoli OMEGAMON XE on z/VM and Linux Troubleshooting Guide.</i>
Diagnosing problems with Tivoli Enterprise Portal	Installation troubleshooting	<i>IBM Tivoli Monitoring Troubleshooting Guide</i>
Diagnosing problems with OMEGAMON XE on z/VM and Linux	General troubleshooting tips for the platform and the monitoring agent	<i>IBM Tivoli OMEGAMON XE on z/VM and Linux Troubleshooting Guide</i>

Using

Table 30 lists the location of information required to use the components of the IBM Tivoli Monitoring and the OMEGAMON XE on z/VM and Linux monitoring agent.

Table 30. Use information for the IBM Tivoli Monitoring and monitoring agents

Component	Information contents	Information location
Using IBM Tivoli Monitoring	From an end-user perspective, using IBM Tivoli Monitoring is a function of using the interface. The use publication for IBM Tivoli Monitoring is the Tivoli Enterprise Portal user's guide. This publication covers the following topics: <ul style="list-style-type: none"> • Overview of IBM Tivoli Monitoring and agents • Starting Tivoli Enterprise Portal and understanding the interface • Tutorial lesson on real-time and event-based monitoring • Editing and customizing workspaces • Using and customizing table and chart views • Building a custom Navigator view • Event views, specifically the message log view, the event console view, and the graphic view • Situations for event-based monitoring, including starting, stopping, customizing, and deleting situations • Automation policies, including creating and maintaining them using the workflows window • Using the terminal view • Functions that can be used as formulas 	<i>IBM Tivoli Monitoring Enterprise Portal User's Guide</i> and Tivoli Enterprise Portal Online Help
Using OMEGAMON XE on z/VM and Linux	<ul style="list-style-type: none"> • Understanding of the Tivoli Enterprise Portal function specific to OMEGAMON XE on z/VM and Linux • Interface overview • Using attributes and workspaces • Using situations and situation events • Using the predefined OMEGAMON XE on z/VM and Linux workspaces • Using the predefined OMEGAMON XE on z/VM and Linux situations • Mapping of workspaces to attribute tables 	<i>IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide</i> and OMEGAMON XE on z/VM and Linux Online Help System
Using Tivoli Enterprise Portal or one of the IBM Tivoli Monitoring monitoring agents	Information about the workspaces, attributes, situations, and expert advice that make up the monitoring agent interface on Tivoli Enterprise Portal.	Tivoli Enterprise Portal Online Help System or OMEGAMON XE on z/VM and Linux Online Help System

Table 30. Use information for the IBM Tivoli Monitoring and monitoring agents (continued)

Component	Information contents	Information location
Using OMEGAMON XE on z/VM and Linux historical data	Information about filters and queries <ul style="list-style-type: none"> • Short-term historical data • Long-term historical data • Monitoring scenarios 	<i>IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide</i>

Support for problem solving

If you have a problem with your IBM software, you want to resolve it quickly. This section describes the following options for obtaining support for IBM software products:

- “Using IBM Support Assistant”
- “Obtaining fixes”
- “Receiving weekly support updates” on page 102
- “Contacting IBM Software Support” on page 102

Using IBM Support Assistant

The IBM Support Assistant is a free, stand-alone application that you can install on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products you use.

The IBM Support Assistant saves you the time it takes to search the product, support, and educational resources. The IBM Support Assistant helps you gather support information when you need to open a problem management record (PMR), which you can then use to track the problem.

The product-specific plug-in modules provide you with the following resources:

- Support links
- Education links
- Ability to submit problem management reports

For more information, and to download the IBM Support Assistant, see <http://www.ibm.com/software/support/isa>. After you download and install the IBM Support Assistant, follow these steps to install the plug-in for your Tivoli product:

1. Start the IBM Support Assistant application.
2. Select **Updater** on the Welcome page.
3. Select **New Properties and Tools** or select the **New Plug-ins** tab (depending on the version of IBM Support Assistant installed).
4. Under **Tivoli**, select your product, and then click **Install**. Be sure to read the license and description. If your product is not included on the list under **Tivoli**, no plug-in is available yet for the product.
5. Read the license and description, and click **I agree**.
6. Restart the IBM Support Assistant.

Obtaining fixes

A product fix might be available to resolve your problem. To determine which fixes are available for your Tivoli software product, follow these steps:

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**.
If you click **Go**, the **Search within all of Tivoli support** section is displayed. If you don't click **Go**, you see the **Select a product** section.
3. Select your product and click **Go**.
4. 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 **Search**.

You can find information specific to the OMEGAMON XE on z/VM and Linux monitoring agent at the following Web address:

<http://www-306.ibm.com/software/sysmgmt/products/support/IBMTivoliOMEGAMONonzVMLinux.html>.

The support page for this monitoring agent displays.

For more information about the types of fixes that are available, see the *IBM Software Support Handbook* at <http://techsupport.services.ibm.com/guides/handbook.html>.

Receiving weekly support updates

To receive weekly e-mail notifications about fixes and other software support news, follow these steps:

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. For example, IBM Tivoli OMEGAMON XE on z/VM and Linux.
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**.

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 (see “Using IBM Support Assistant” on page 101).

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 Web site 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 Web site 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 Web site at <https://techsupport.services.ibm.com/ssr/login>.
- For customers with IBMLink, CATIA, Linux, OS/390, iSeries®, pSeries®, zSeries, and other support agreements, go to the IBM Support Line Web site 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 Web site 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://techsupport.services.ibm.com/guides/contacts.html> and click the name of your geographic region for phone numbers of people who provide support for your location.

To contact IBM Software support, follow these steps:

1. “Determining the business impact”
2. “Describing problems and gathering information”
3. “Submitting problems” on page 104

Determining the business impact

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

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.

Documentation library

This appendix contains information about the publications related to OMEGAMON XE on z/VM and Linux and to IBM Tivoli Monitoring and the commonly shared components of Tivoli Management Services. These publications are listed in the following categories:

- “OMEGAMON XE on z/VM and Linux library”
- “IBM Tivoli Monitoring publications” on page 106
- “Related publications” on page 107

See *IBM Tivoli Monitoring and OMEGAMON XE Products: Documentation Guide*, SC23-8816, for information about accessing and using the publications. You can find the Documentation Guide, under **Monitoring**, in the IBM Tivoli Monitoring and OMEGAMON XE Information Center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/>.

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 information centers** on the Welcome page for the product.

OMEGAMON XE on z/VM and Linux library

The following publications comprise the OMEGAMON XE on z/VM and Linux library:

- *IBM Tivoli OMEGAMON XE on z/VM and Linux Program Directory*, GI11-4135
Provides hardware and software prerequisites for the OMEGAMON XE on z/VM and Linux installation and instructions for the VMSES/E part of the installation.
- *IBM Tivoli OMEGAMON XE on z/VM and Linux Planning and Configuration Guide*, SC27-2837
Provides information for installing and configuring OMEGAMON XE on z/VM and Linux.
- *IBM Tivoli OMEGAMON XE on z/VM and Linux User's Guide*, SC27-2836
Introduces the features and describes the workspaces, attributes, and predefined situations for the OMEGAMON XE on z/VM and Linux product. Supplements the online help provided with this product by including product-specific monitoring scenarios.
- *IBM Tivoli OMEGAMON XE on z/VM and Linux Troubleshooting Guide*, GC27-2838
Provides information and messages to assist you in troubleshooting problems with the OMEGAMON XE on z/VM and Linux monitoring agent.
- *IBM Tivoli OMEGAMON XE on z/VM and Linux Readme* , GI11-9442
Contains late-breaking information about the OMEGAMON XE on z/VM and Linux product limitations and workarounds.
- *IBM Tivoli OMEGAMON XE on z/VM and Linux Quick Start Guide* , GI11-9441
Provides overview information to get you started installing and configuring OMEGAMON XE on z/VM and Linux.

You can access the OMEGAMON XE on z/VM and Linux library at the following Web address:

http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.omegamon_xezvm.doc_4.2/welcome.htm

You can find additional information at the following Web address: <http://www.ibm.com/software/sysmgmt/products/support>

Select **IBM Tivoli OMEGAMON XE on z/VM and Linux** from the **Support for specific Tivoli products** drop-down list. You will be taken to the support page for this monitoring agent.

OMEGAMON XE on z/VM and Linux online help

The online help aids operators in understanding and using the provided data, attributes, and situations to monitor performance and availability in the context of this product.

Note: This set of information is nested within the Tivoli Enterprise Portal help, which describes global features that are useful in helping you to use OMEGAMON XE on z/VM and Linux. For more information, see “Tivoli Enterprise Portal help system” on page 107.

IBM Tivoli Monitoring publications

To use the information in this guide effectively, you must have some prerequisite knowledge about IBM Tivoli Monitoring (also called IBM Tivoli Management Services) and the Tivoli Enterprise Portal interface, which you can obtain from the following guides. The publications are available in the IBM Tivoli Monitoring Information Center at the following Web address:

<http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/>

- *IBM Tivoli Monitoring Installation and Setup Guide*, GC32-9407
Provides information on installing and setting up the Tivoli Enterprise Monitoring Server and the Tivoli Enterprise Portal Server and client. It also describes how to install the distributed monitoring agents.
- *Configuring IBM Tivoli Enterprise Monitoring Server on z/OS*, SC27-2313
Describes how to configure and customize the Tivoli Enterprise Monitoring Server on z/OS. The publication also contains planning information and information about setting up security on your Tivoli Enterprise Monitoring Server.
- *IBM Tivoli Monitoring Administrator's Guide*, SC32-9408
Describes the support tasks and functions required for the Tivoli Enterprise Portal Server and clients, including Tivoli Enterprise Portal user administration.
- *IBM Tivoli Monitoring Enterprise Portal User's Guide*, SC32-9409
Describes how to use the Tivoli Enterprise Portal interface. It includes a tutorial about monitoring that covers workspaces, navigation, views, and responding to alerts. Different types of views and situations for event-based monitoring are also included, as well as information on automation policies.
- *Exploring IBM Tivoli Monitoring*, SC32-1803
Provides a series of exercises that help users explore ITM Tivoli Monitoring.
- *IBM Tivoli Universal Agent API and Command Programming Reference Guide*, SC32-9461
Contains the procedures for implementing the IBM Tivoli Universal Agent application programming interfaces (APIs) and provides descriptions, syntax, and return status codes for the API calls and command-line interface commands.
- *IBM Tivoli Monitoring Troubleshooting Guide*, GC32-9458
Provides information and messages to assist users with troubleshooting problems with the software.
- *IBM Tivoli Monitoring: Upgrading from Tivoli Distributed Monitoring*, GC32-9459
Provides information on how to upgrade from Tivoli Distributed Monitoring.
- *IBM Tivoli Universal Agent User's Guide*, SC32-9459
Introduces you to the IBM Tivoli Universal Agent, an agent of IBM Tivoli Monitoring. The IBM Tivoli Universal Agent enables you to use the monitoring and automation capabilities of IBM Tivoli Monitoring to monitor any type of data you collect.
- *Introducing IBM Tivoli Monitoring*, GI11-4071
Gives a basic introduction to the features of IBM Tivoli Monitoring.
- *IBM Tivoli OMEGAMON XE and Tivoli Management Services on z/OS Upgrade Guide*, GC32-1980

Provides instructions for performing a staged upgrade for OMEGAMON XE V4.1.0 monitoring agents, discusses the basic upgrade requirements and the effects of running a mixed environment for globalization and data presentation. This guide also lists the required sequence of steps for performing a staged upgrade in two scenarios.

Additionally, the following user's guide provides information about the IBM Tivoli Monitoring for Linux OS agent. The OMEGAMON XE on z/VM and Linux monitoring agent contains dynamic links to workspaces in the IBM Tivoli Monitoring for Linux OS agent. See "Dynamic linking to cross-product workspaces" on page 9 for details.

- *IBM Tivoli Monitoring: Agent for Linux OS User's Guide*, SC32-9447
Provides information about using the IBM Tivoli Monitoring for Linux OS agent. This includes descriptions of the workspaces, the situations, and the attributes provided by this agent.

Tivoli Enterprise Portal help system

The Tivoli Enterprise Portal help system provides context-sensitive reference information regarding IBM Tivoli Monitoring. The help system also documents features and customization options, such as the Link Wizard that helps you to modify the default linkages provided in the pop-up menus of rows in a workspace.

Related publications

You can access the entire library for the IBM z/VM Operating System at the following Web address:

<http://www.vm.ibm.com/library/>

The following z/VM publications are relevant to this monitoring agent:

- *z/VM: Performance Toolkit Guide*, SC24-6156
This publication provides, in concert with the SC24-6157, contains all information required for using the Performance Toolkit for VM (previously known as the VM/ESA[®] Full Screen Operator Console and Graphical Real Time Monitor, FCON/ESA, or just FCON).
- *z/VM: Performance Toolkit Reference*, SC24-6157
This publication, in concert with the *z/VM: Performance Toolkit Guide*, contains all information required for using the Performance Toolkit for VM (previously known as the VM/ESA Full Screen Operator Console and Graphical Real Time Monitor, FCON/ESA, or just FCON).
You can also access information on the Performance Toolkit and on the latest enhancements to the Performance Toolkit as they pertain to this monitoring agent at the following Web address:
<http://www.vm.ibm.com/related/perfkit/pksegout.html>
- *z/VM: CP Commands and Utilities Reference*, SC24-6081
Lists and describes IBM z/VM Control Program (CP) commands and utilities for users of every privilege class.
- *z/VM: CMS Commands and Utilities Reference*, SC24-6073
Provides reference information about Conversational Monitor System (CMS) commands and utilities for IBM z/VM.
- *Device Drivers, Features and Commands*, SC33-8281
Provides information about the device drivers available to Linux for the control of zSeries[®] and S/390[®] devices and attachments with the kernel 2.6 (April 2004 stream). It also provides information on commands and parameters relevant to configuring Linux for zSeries and S/390[®].
- *Saved Segments Planning and Administration*, SC24-6116
Provides general information on Discontinuous Saved Segments (DCSS).
- *Directory Maintenance VM/ESA V1R5.0 Command Reference*, SC20-1839

Provides reference information for the z/VM Directory Maintenance Facility (DirMaint™) Function Level 510, for use on IBM z/VM Version 5.

You may also want to review the IBM Redbooks for z/VM, as well as those listed below, at the following Web address:

<http://www.vm.ibm.com/pubs/redbooks/>

- *Running Linux on IBM System z9® and zSeries under z/VM*, SG24-6311
- *z/VM and Linux on IBM System z: The Virtualization Cookbook for SLES9*, SG24-6695
- *Linux on IBM eServer zSeries and S/390: Performance Measurement and Tuning*,

You can find links to hints and tips about z/VM performance at the following Web site:

<http://vm.ibm.com/perf>

Additionally, the following Web site contains links to a collection of observations from the Linux Performance Team for Linux on System z:

<http://www.vm.ibm.com/developerworks/linux/linux390/perf/index.html>

Accessing terminology online

The *Tivoli Software Glossary* includes definitions for many of the technical terms related to Tivoli software. The *Tivoli Software Glossary* is available at the following Tivoli software library Web site:

<http://publib.boulder.ibm.com/tividd/glossary/tivologlossarymst.htm>

The IBM Terminology Web site consolidates the terminology from IBM product libraries in one convenient location. You can access the Terminology Web site at the following Web address:

<http://www.ibm.com/ibm/terminology>

Accessing publications online

The documentation CD contains the publications that are in the product library. The format of the publications is PDF, HTML, or both. Refer to the readme file on the CD for instructions on how to access the documentation.

IBM posts publications for this and all other Tivoli products, as they become available and whenever they are updated, to the Tivoli Information Center Web site at the following Web address:

<http://publib.boulder.ibm.com/tividd/td/link/tdprodlist.html>.

In the Tivoli Software Information Center window, click the letter that matches the first letter of your product name to access your product library. For example, click **M** to access the IBM Tivoli Monitoring library or click **O** to access the IBM Tivoli OMEGAMON library.

Note: If you print PDF documents on other than letter-sized paper, set the option in the **File → Print** window that allows Adobe® Reader to print letter-sized pages on your local paper.

Ordering publications

You can order many Tivoli publications online at the following Web site:

<http://www.elink.ibm.link.ibm.com/publications/servlet/pbi.wss>

You can also order by telephone by calling one of these numbers:

- In the United States: 800-879-2755
- In Canada: 800-426-4968

In other countries, contact your software account representative to order Tivoli publications. To locate the telephone number of your local representative, perform the following steps:

1. Go to <http://www.elink.ibm.link.ibm.com/publications/servlet/pbi.wss>.
2. Select your country from the list and click **Go**.
3. Click **About this site** in the main panel to see an information page that includes the telephone number of your local representative.

Accessibility

Accessibility features help users with a physical disability, such as restricted mobility or limited vision, to use software products successfully. With this product, you can use assistive technologies to hear and navigate the interface. You can also use the keyboard instead of the mouse to operate all features of the graphical user interface.

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:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, 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. _enter the year or years_. All rights reserved.

If you are viewing this information in softcopy form, the photographs and color illustrations might not display.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol ([®] or [™]), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe[®], Acrobat, Portable Document Format (PDF), and PostScript[®] are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

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.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Glossary

A

adding application support. Before you can use a monitoring agent, the Tivoli Enterprise Monitoring Server to which it reports must be initialized with application data. This step adds product-provided situations, templates, and other sample data to the Tivoli Enterprise Monitoring Server Enterprise Information Base (EIB) tables.

application data. The product-provided situations, templates, and other sample data to the Enterprise Information Base (EIB) tables of the Tivoli Enterprise Monitoring Server.

attribute. A system or application element being monitored by the monitoring agent, such as Disk Name and Disk Read/Writes Per Second. An attribute can also be a field in an ODBC-compliant database.

attribute table. A set of related attributes that can be combined in a data view or a situation. When you open the view or start the situation, Tivoli Enterprise Portal retrieves data samples of the selected attributes. Each type of agent has a set of attribute groups.

authorized virtual machine. In z/VM, a GCS virtual machine associated with an authorized user ID.

C

client/server. In communications, the model of interaction in distributed data processing in which a program at one site sends a request to a program at another site and awaits a response. The requesting program is called a client; the answering program is called a server.

command. A request from a user to the z/VM system to perform a particular operation. A CMS command can also be the name of a CMS file with a file type of EXEC or MODULE, in which a sequence of operations are defined. See also **Conversational Monitor System (CMS)**.

Configuration Tool. Tool used to install the previous Candle products and to configure some of the Tivoli Monitoring Agent zSeries products (which are now installed using the System Modification Program/Extended (SMP/E) tool).

Control Program (CP). A component of the z/VM system that manages the resources of a single computer so that multiple computing systems seem to exist. Each apparent system, or virtual machine, is the functional equivalent of the real computer, and CP simulates the real machine architecture in the virtual machine.

Conversational Monitor System (CMS). A component of z/VM that runs in a virtual machine and provides both the interactive z/VM end-user interface and the general z/VM application programming interface. CMS runs only under the control of the z/VM Control Program (CP).

D

device driver. In z/VM: (1) A file that contains the code needed to use an attached device. (2) A program that enables a computer to communicate with a specific peripheral device, such as a printer or a CD drive. (3) A collection of subroutines that control the interface between I/O device adapters and the processor. (4) In CMS Pipelines, a stage that reads data from or writes data to I/O and storage devices, host environments (including CP, CMS, and XEDIT), and REXX and EXEC 2 variables.

direct access storage device (DASD). In z/VM, a mass storage medium in which the data access time is effectively independent of the data location.

discontiguous saved segment (DCSS). In z/VM, a saved segment that begins and ends on a megabyte boundary and is not a segment space or a member of a segment space. A DCSS can contain logical saved segments. See also **segment space**.

E

expert advice. A description within the Situation Editor of each situation provided with a monitoring agent to help you quickly gather and interpret data.

G

guest operating system. An operating system, such as Linux or z/OS, running in a virtual machine managed by the z/VM Control Program (CP).

guest virtual machine. In z/VM, a virtual machine in which an operating system is running.

guest virtual storage. In z/VM, the storage that appears to the operating system running in a virtual machine.

H

HiperSockets. (1) A hardware channel that provides high-speed TCP/IP communication between logical partitions (LPARs) on the same IBM zSeries server. It uses an adaptation of the queued direct I/O (QDIO) architecture. (2) The virtualization of the HiperSockets

channel in z/VM, which provides high-speed communication between guest virtual machines.

historical data management. The procedures applied to short-term binary history files that perform roll off to either a data warehouse or to delimited text files and delete entries in the short-term history files over 24 hours old to make room for new entries.

host. The z/VM Control Program (CP) in its capacity as manager of a virtual machine in which another operating system is running.

hub Tivoli Enterprise Monitoring Server. The Tivoli Enterprise Monitoring Server that has been elected to act as the focal point to which all Tivoli Enterprise Portal Servers connect.

I

IBM Tivoli Monitoring. A client-server implementation comprising a Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Tivoli Enterprise Portal client, and monitoring agents that collect and distribute data to the Tivoli Enterprise Monitoring Server.

inter-user communication vehicle (IUCV). In z/VM, a CP interface for passing data between virtual machines or between CP and a virtual machine.

M

member saved segment. In z/VM, a saved segment that begins and ends on a page boundary and belongs to up to 64 segment spaces. A member saved segment is accessed by its own name or by the name of a segment space to which it belongs. A member saved segment may contain logical saved segments. See discontinuous saved segment (DCSS) and segment space.

monitoring agent. The agent process probes a managed system for data and sends information back to Tivoli Enterprise Portal formatted into table and chart views.

monitoring interval. A specified time, scalable to seconds, minutes, hours, or days, for how often the Tivoli Enterprise Monitoring Server checks to see if a situation has become true. The minimum monitoring interval is 30 seconds; the default is 15 minutes.

L

logical partition (LPAR). In z/VM, a subset of the processor hardware that is defined to support the operation of a system control program (SCP).

N

Navigator. The left pane of the Tivoli Enterprise Portal window. The Navigator Physical view shows your network enterprise as a physical hierarchy of systems grouped by platform. OMEGAMON DE users can also create other views to create logical hierarchies grouped as you specify, such as by department or function.

O

OMEGAMON Web Services. An open standards-based interface to IBM Tivoli Monitoring using SOAP requests. Any Tivoli Monitoring Agent can be dynamically queried, so performance and availability data can be processed by other applications.

P

paging. In z/VM, transferring pages between real storage and external page storage.

Performance Toolkit for VM. An optional feature of z/VM that gathers, analyzes, and displays VM performance data. It can also process Linux performance data obtained from the Resource Management Facility.

privilege class. In z/VM, the authorization required to use a subset of the CP commands and DIAGNOSE codes and certain CP system functions. The privilege class corresponds to a type of user. Each CP command and DIAGNOSE code belongs to one or more privilege classes. A user is assigned one or more privilege classes in the z/VM directory entry for that virtual machine.

R

remote Tivoli Enterprise Monitoring Server. The Tivoli Enterprise Monitoring Server that passes its collected data to the hub Tivoli Enterprise Monitoring Server to be made available to clients, creating an enterprise-wide view.

S

seeding. See **adding application support**.

shared segment. See **discontinuous saved segment (DCSS)**.

situation. A set of conditions that, when met, creates an event. A condition consists of an attribute, an operator such as greater than or equal to, and a value. It can be read as, "If - system condition - compared to - value - is true." An example of a situation is: IF - CPU usage - GT - 90% - TRUE. The expression "CPU usage GT 90%" is the situation condition.

T

Take Action. A command window on Tivoli Enterprise Portal from which you can enter your command or choose from a list of predefined commands. It also has a list of systems on which to effect the command.

target libraries. SMP/E-controlled libraries that contain the data from the distribution media

threshold. A level set in the system at which a message is sent or an error-handling program is called. For example, in a user auxiliary storage pool, the user can set the threshold level in the system values, and the system notifies the system operator when that level is reached.

Tivoli Enterprise Monitoring Server. The IBM Tivoli Monitoring component that retrieves data from the monitoring agents and delivers data to the Tivoli Enterprise Portal Server, sends alerts to the Tivoli Enterprise Portal Server when conditions specified in situations are met, receives commands from the Tivoli Enterprise Portal and passes them to the appropriate monitoring agents, and (optionally) provides a repository for short-term historical data. This component can be installed on z/OS, Windows, and some UNIX operating systems.

Tivoli Enterprise Portal Server. The IBM Tivoli Monitoring server you log on to. The Tivoli Enterprise Portal Server connects to the hub Tivoli Enterprise Monitoring Server. It enables retrieval, manipulation and analysis of data from IBM Tivoli Monitoring managed systems.

V

view. A windowpane, or frame, in a workspace. It might contain data from an agent in a chart or table, or it might contain a terminal session or browser, for example. A view can be split into two separate, autonomous views.

W

working set. The estimated number of pages of real storage that a virtual machine needs to run.

workspace. The viewing area of the Tivoli Enterprise Portal window, excluding the Navigator. Each workspace comprises one or more views. Every Navigator item has its own default workspace and might have multiple workspaces.

Index

A

- accessibility 109
- Action tab
 - enabling Take Action commands 53
- administrative tasks 98
- audience
 - expertise xi
 - responsibilities xi

C

- communication protocols
 - hub monitoring server on z/OS 86
 - monitoring agent 81
 - portal desktop client 92
 - portal server 83
 - Tivoli Enterprise Monitoring Server
 - distributed 85
 - Tivoli Enterprise Portal 92
 - Tivoli Enterprise Portal Server 83
- communication protocols between components
 - worksheets 93
- components
 - IBM Tivoli Monitoring 16
- configuration 97
 - completing 63
 - configuration steps 40
 - configuring TCP/IP on z/VM 40
 - configuring the device driver on the Linux guest 47
 - enabling Take Action commands 53
 - enabling the collection of data on z/VM 40
 - enabling the collection of Linux data 49
 - enabling the CP Monitor domains 40
 - estimating the size of the DCSS 42
 - guidelines for issuing Take Action commands 56
 - installation prerequisites 37
 - Linux guest requirements for Take Action commands 54
 - monitoring agent 80
 - overall 79
 - planning 15
 - required order of tasks for viewing data at the monitoring agent 38
 - running out of memory in the DCSS 46
 - setting the environment variables for Take Action commands 54
 - starting data collection for User AppIData at the Linux guest 50
 - starting data collection for User AppIData at the Linux guest automatically 51
 - Take Action commands to be excluded from running 57
 - Take Action commands you can issue 56
 - Tivoli Enterprise Monitoring Server on distributed systems 31
 - Tivoli Enterprise Monitoring Server on z/OS 31
 - z/VM requirements for Take Action commands 53

- customer support
 - See Software Support

D

- data warehouse
 - proxy agent 19
 - Warehouse Proxy planning 19
 - Warehouse Summarization and Pruning agent 19
- DCSS
 - adding the DCSS to the Linux guest 49
 - calculating the size of the DCSS 42
 - configuring the device driver on the Linux guest 47
 - defining 45
 - defining the guest storage with storage gaps 47
 - determining the start and end addresses of the PERFOUT DCSS 47
 - estimating the size of the DCSS 42
 - extending the Linux address range 48
 - helpful tips for defining your own segment 45
 - loading the DCSS at startup time 49
 - loading the device driver 48
 - naming scheme 49
 - running out of memory 46
 - using FCXSEGSZ 42
- defects, limitations, and workarounds 96
- defining
 - discontiguous saved segment (DCSS) 45
- defining user IDs
 - for Linux guests 61
 - in IBM Tivoli Monitoring 61
 - security 61
- diagnosis 99
- discontiguous saved segment (DCSS) 45
- dynamic linking to cross-product workspaces 9
- dynamic workspace linking
 - enabling dynamic workspace linking 52

E

- enabling Take Action commands
 - loading the vmcp device driver 56
 - verifying that sudo is available 56
- enabling the collection of Linux data 49
- enabling the CP Monitor domains
 - EVENT APPLDATA command 41
 - SAMPLE APPLDATA command 41

F

- finding the information you need 95
- fixes, obtaining 101

G

- guide
 - what this guide contains xi

H

- hardware
 - prerequisites 25
 - required 25
- historical data collection 7, 21, 66
 - long-term 66
 - setting up 32
 - short-term 66
 - types of data to collect 67
- hub Tivoli Enterprise Monitoring Server
 - defined 18

I

- IBM Redbooks 101
- IBM Support Assistant 101
- IBM Tivoli Monitoring
 - architecture 16
 - communication protocols 93
 - components 16
 - features 6
 - upgrading 33
- IBM Tivoli Monitoring publications 106
- IBM Tivoli OMEGAMON XE
 - zSeries products 10
- IBM Tivoli OMEGAMON XE on z/VM and Linux
 - distribution media 27
 - product package 26
- installation 97
 - distributed installation 27
 - installation flow 29
 - preparation 21
 - z/VM installation 28
- installation flow 29
- installing Language Support 57
- interoperability with other products 13

L

- Linux operating system
 - supported versions 24
- log files
 - data warehouse 75
 - definitions of variables for RAS1 logs 74
 - locations 73
 - monitoring server on Windows or UNIX 75
 - portal client 75
 - portal server 75

M

- manuals
 - see documentation library 105
 - see publications 108
- monitoring agent
 - communication protocols worksheet 81
- monitoring server
 - hub 18
 - remote 18
 - worksheet, distributed 84

- monitoring servers
 - hub and remote 17

N

- Navigator views 7

O

- OMEGAMON XE on z/VM and Linux 3
 - architecture overview 10
 - overview 3
 - packaging 25
 - planning worksheets 79
 - sample workspace 4
- OMEGAMON XE on z/VM and Linux monitoring agent 9
- OMEGAMON XE on z/VM and Linux product publications 105
- online publications
 - accessing 108
- ordering publications 109

P

- packaging
 - tape formats 25
- performance considerations 65
 - disk capacity planning for historical data 68
 - historical data collection 66
 - real-time data collection 65
 - situations 69
 - systems to monitor 66
 - workspace design 70
- Performance Toolkit
 - guides and Web site 107
 - role in this monitoring agent 3
- planning 30, 95
 - configuration 15
 - worksheets 15, 30
- planning the deployment 30
- planning worksheets 79
- portal client
 - deciding between Windows and Linux 16
- portal desktop client
 - Linux configuration worksheet 91
 - Windows configuration worksheet 91
- portal desktop client on Windows
 - communication protocols 92
- portal server
 - communication protocols 83
 - deciding between Windows and Linux 16
 - Linux configuration worksheet 82
 - Windows configuration worksheet 82
- preparing your z/VM environment 30
- prerequisites
 - hardware 21
 - software 21
 - OMEGAMON XE on z/VM and Linux 24
 - Tivoli Data Warehouse 21

- prerequisites *(continued)*
 - software *(continued)*
 - Tivoli Enterprise Monitoring Server on distributed 21
 - Tivoli Enterprise Portal 21
 - Tivoli Enterprise Portal Server 21
 - Tivoli Monitoring Agent for Linux OS 25
 - z/VM operating system software requirements 22
- problem determination
 - describing problems 103
 - determining business impact 103
 - submitting problems 104
- problem resolution 101
- product documentation
 - distribution media 28
- product packaging 25
- protocols, communication
 - See communication protocols
- proxy agent 19
- publications 105
 - accessing online 108
 - IBM Tivoli Monitoring 106
 - OMEGAMON XE on z/VM and Linux product 105
 - ordering 109
 - related 107
 - see documentation library 105
 - see publications 108

Q

- queries 8

R

- real-time data collection 65
- Red Hat Enterprise Linux operating system
 - supported versions 24
- Redbooks 101
- remote Tivoli Enterprise Monitoring Server
 - defined 18
- required order of tasks for viewing data at the monitoring agent 38
- running the Linux guest as a non-root user
 - on Red Hat Enterprise Linux 5 55
- running under Linux as a non-root user
 - on SUSE Linux Enterprise Server 9 for zSeries 55
 - Take Action commands 54
- running under Linux guest as a non-root user
 - on SUSE Linux Enterprise Server 10 for zSeries 55

S

- security
 - configuring 31
 - defining user IDs 61
 - defining user IDs for the Command Processor 61
 - enabling 61
 - validation 62
- serviceability 12
 - log file locations 73
- Situation editor 8

- situations
 - autostarting to improve performance 69
 - defining 69
 - grouping to improve performance 69
 - running 69
- Situations
 - using the Action tab 53
- software
 - prerequisites 21
 - required 21
- Software Support
 - contacting 102
 - describing problems 103
 - determining business impact 103
 - overview 101
 - receiving weekly updates 102
 - submitting problems 104
- standards supported 12
- starting the monitoring agent 58
- stopping the monitoring agent 59
- storage considerations 65
- support assistant 101

T

- Take Action commands
 - commands to exclude 57
 - commands you can issue 56
 - enabling Take Action commands 53
 - environment variables 54
 - guidelines for issuing 56
 - Linux guest requirements 54
 - running under Linux as a non-root user 54
 - verifying that commands are being issued 56
 - z/VM requirements 53
- task map 95
 - administration 98
 - configuring 97
 - diagnosis 99
 - installing 97
 - planning 95
 - temporary defects 96
 - tuning 98
 - uninstalling 97
 - upgrading 96
 - using 99
- Tivoli Data Warehouse 19
- Tivoli Enterprise Console 19
- Tivoli Enterprise Monitoring Server
 - hub 18
 - worksheet, distributed 84
- Tivoli Enterprise Monitoring Servers 17
- Tivoli Enterprise Portal
 - communication protocols 92
- Tivoli Enterprise Portal client
 - deciding between Windows and Linux 16
- Tivoli Enterprise Portal desktop client
 - configuration worksheet 91
- Tivoli Enterprise Portal Server
 - communication protocols 83
 - configuration worksheet 82

- Tivoli Enterprise Portal Server *(continued)*
 - deciding between Windows and Linux 16
- Tivoli Management Services
 - components 6
- Tivoli Monitoring Agent for Linux OS
 - software
 - prerequisite 25
- Tivoli software information center 108
- troubleshooting 99
- tuning 98

U

- Universal Agent support 8
- upgrading 96
 - existing IBM Tivoli Monitoring environment 33
- using 99

V

- verification 31
- version of
 - IBM Tivoli Monitoring 33
- views
 - browser view 7
 - message log view 7
 - notepad view 7
 - table views 7
 - take action views 7
 - terminal view 7
- VMSES/E 29

W

- Warehouse Summarization and Pruning agent 19
- worksheets
 - hub monitoring server on a distributed system 84
 - hub monitoring server on z/OS 86
 - hub Tivoli Enterprise Monitoring Server on a distributed system 84
 - monitoring agent 80, 81
 - monitoring server, distributed 85
 - overall configuration 79
 - portal desktop client configuration 91
 - portal server configuration 82
 - Tivoli Enterprise Monitoring Server, distributed 85
 - Tivoli Enterprise Portal desktop client configuration 91
 - Tivoli Enterprise Portal Server configuration 82
- workspaces
 - designing 70
 - auto-refresh rate 71
 - number of attributes retrieved 70
 - number of rows retrieved 70
 - queries to multiple views 71
 - views 7

Z

- z/VM environment
 - preparing 30
- z/VM operating system software requirements 22



Printed in USA

SC27-2837-00

