IBM Tivoli Monitoring
Version 6.2.3 Fix Pack 1

UNIX OS Agent User's Guide

SC32-9446-05

IBM Tivoli Monitoring
Version 6.2.3 Fix Pack 1

UNIX OS Agent User's Guide

SC32-9446-05

Note
Before using this information and the product it supports, read the information in "Notices" on page 333.

This edition applies to version 6.2.3 Fix Pack 1 of IBM Tivoli Monitoring: UNIX OS Agent (product number 5724-C04) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Tables vii

Chapter 1. Overview of the Monitoring Agent for UNIX OS 1

IBM Tivoli Monitoring overview 1
Features of the Monitoring Agent for UNIX OS . . . 1
New in this release 2
Monitoring Agent for UNIX OS components 5
User interface options 6

Chapter 2. Requirements for the monitoring agent 7

Enabling the Monitoring Agent for UNIX OS to run as a nonroot user 9
Securing your IBM Tivoli Monitoring installation 9
Setting overall file ownership and permissions for nonroot users 9
Setting kuxagent binary permissions 9
Limitations of starting the agent as a nonroot user 10
Setting up the Monitoring Agent for UNIX OS in a cluster environment 11
Using Agent Management Services 11
Filtering capabilities on the names of processes . . 12

Chapter 3. Workspaces reference . . . 13

About workspaces 13
More information about workspaces 13
Predefined workspaces 13
Agent Management Services workspace 15
Agents' Management Log workspace 15
AIX Devices Status workspace 16
AIX LPAR Information workspace 16
AIX WPAR Summary workspace 16
AIX WPAR CPU Details workspace 17
AIX WPAR Details workspace 17
AIX WPAR Memory Details workspace 17
AIX WPAR Network and File System workspace 18
All Files workspace 18
All Processes workspace 18
Disk Usage workspace 18
Disk Usage Details workspace 19
Disk Utilization for Mount Point workspace . . 19
Enterprise UNIX System Summary workspace . . 19
File Information workspace 20
Historical Summarized Availability workspace . . 20
Historical Summarized Availability Daily workspace 20
Historical Summarized Availability Hourly workspace 20
Historical Summarized Availability Weekly workspace 21
Historical Summarized Capacity workspace . . 21
Historical Summarized Capacity Daily workspace 22
Historical Summarized Capacity Hourly workspace 22
Historical Summarized Capacity Weekly workspace 23
Historical Summarized Performance workspace 24
Historical Summarized Performance Daily workspace 24
Historical Summarized Performance Hourly workspace 25
Historical Summarized Performance Weekly workspace 26
Network workspace 26
NFS Activity workspace 27
Process workspace 28
RPC Performance workspace 29
Solaris System CPU Workload workspace . . . 29
Solaris Zone Processes workspace 29
Solaris Zones workspace 29
System Details workspace 30
System Information workspace 30
Top CPU-Memory %-VSize Details workspace . . 31
UNIX workspace 31
UNIX Detail workspace 31
Users workspace 32

Chapter 4. Attributes reference 33

About attributes 33
More information about attributes 34
Attribute groups and attributes 34
Agent Availability Management Status attributes . . 36
Agent Active Runtime Status attributes 37
AIX AMS attributes 38
AIX Defined Users attributes 39
AIX Devices attributes 39
AIX LPAR attributes 40
AIX WPAR CPU attributes 41
AIX WPAR File System attributes 42
AIX WPAR Information attributes 43
AIX WPAR Network attributes 44
AIX WPAR Physical Memory attributes 44
Alerts Table attributes 45
All Users attributes 46
Configuration Information attributes 47
Disk attributes 48
Disk Performance attributes 52
File Comparison Group attributes 54
File Information attributes 55
File Pattern attributes 58
Group attributes 58
IP Address attributes 58
Machine Information attributes 59
Network attributes 60
NFS / RPC Statistics attributes 66
Ping attributes 74
Print Queue attributes 75

Process attributes 76
SMP CPU attributes 86
Solaris Zones attributes 90
SP2 System attributes 91
System attributes 99
TCP Statistics Attributes 107
UNIX Memory attributes 107
User attributes 113
Disk capacity planning for historical data 115

Chapter 5. Situations reference 117
About situations 117
More information about situations 118
Predefined situations 118
Agent Management Services node 120
Disk Usage node 120
File Information node 122
Network node 122
NFS Activity workspace 123
Process node 123
System Information node 125

Chapter 6. Take Action commands reference 131
About Take Action commands 131
More information about Take Action commands 131
Predefined Take Action commands 131
AMS Recycle Agent Instance 132
AMS Reset Agent Daily Restart Count 132
AMS Start Agent action 133
AMS Start Agent Instance action 133
AMS Stop Agent action 134
AMS Start Management action 134
AMS Stop Management action 135
Sample_kill_Process action 135

Chapter 7. Policies reference. 137
About policies 137
More information about policies 137
Predefined policies 137
UNIX_CPU_Busy policy 137
UNIX_Disk_Space_Full policy 138
UNIX_Virtual_Memory_High policy 138

Chapter 8. Tivoli Common Reporting for the monitoring agent 139
Utilization Details for Single Resource report . . . 143
Utilization Details for Multiple Resources report 147
Utilization Comparison for Single Resource report 150
Utilization Comparison for Multiple Resources report 152
Utilization Heat Chart for Single Resource report 155
Memory Utilization for Single Resource report . . 159
Memory Utilization for Multiple Resources Comparison report 161
Top Resources Utilization report 164
Top Situations by Status report 168
Enterprise Resources List report 169
Enterprise Daily Utilization Heat Chart report . . 170
Enterprise Summary report 171
Top Resources by Availability 173
Top Resources Utilization Summary Heat Chart report 175
Top Resources by Availability (MTTR/MTBSI) . . 177
Resource Availability Comparison 179
Availability Heat Chart for Single Resource . . . 181
CPU Utilization Comparison for Multiple Resources 183
CPU Utilization for Single Resource 185
Disk Utilization for Single Resource 187
Disk Utilization Comparison for Multiple Resources 190
Situations History report 193
Creating custom queries and reports 195

Chapter 9. Troubleshooting 199
Gathering product information for IBM Software Support 199
Built-in troubleshooting features 199
Problem classification. 200
Trace logging 200
Overview of log file management 200
Examples of trace logging 201
Principal trace log files 201
Setting RAS trace parameters 204
Problems and workarounds 205
Installation and configuration troubleshooting 206
Agent troubleshooting 210
Tivoli Enterprise Portal troubleshooting . . . 215
Troubleshooting for remote deployment . . . 216
Workspace troubleshooting 216
Situation troubleshooting 217
Take Action command troubleshooting 221
Troubleshooting for UNIX 221
Support information 222
Accessing terminology online 222
Accessing publications online 223
Ordering publications 223
Tivoli technical training 223
Tivoli user groups 223

Appendix A. Upgrading for warehouse summarization 225
Tables in the warehouse 225
Effects on summarized attributes 225
Upgrading your warehouse with limited user permissions 226
Types of table changes 227
Table summary 228
Upgrading your warehouse for primary key and tablespace changes 228
Affected attribute groups and supporting scripts 229
Procedures 230

Appendix B. AIX Premium agent attributes 235

Appendix C. IBM Tivoli Enterprise Console event mapping 245

Appendix D. Historical data 275
UNIXALLUSR historical table 275
UNIXAMS historical table 275
UNIXCPU historical table 276
UNIXDEVIC historical table 277
UNIXDISK historical table 277
UNIXDPERF historical table 278
UNIXDUSERS historical table 278
UNIXFILE historical table 278
UNIXGROUP historical table 279
UNIXIPADDR historical table 279
UNIXLPAR historical table 280
UNIXMACHIN historical table 281
UNIXMEM historical table 281
UNIXNET historical table 283
UNIXNFS historical table 283
UNIXOS historical table 285
UNIXPING historical table 288
UNIXPRINTQ historical table 288
UNIXPS historical table 289
UNIXSOLZON historical table. 291
UNIXTCP historical table 291
UNIXUSER historical table 291
UNIXWPARCP historical table 292
UNIXWPARFS historical table 292
UNIXWPARIN historical table 293
UNIXWPARNE historical table 294
UNIXWPARPM historical table 294

Appendix E. Monitoring Agent for UNIX data collection 297

Appendix F. Discovery Library Adapter for the monitoring agent . . . 323
About the DLA. 323
More information about DLAs. 323
UNIX data model class types represented in CDM 323
UNIX class 323
ComputerSystem class 324
IpInterface class 325
IpV4Address class. 325
IpV6Address class. 325
Fqdn class 326
TMSAgent class 326

Appendix G. Documentation library 327
IBM Tivoli Monitoring library 327
Documentation for the base agents 328
Related publications 329
Other sources of documentation 329

Appendix H. Accessibility 331
Navigating the interface using the keyboard . . . 331
Magnifying what is displayed on the screen . . . 331

Notices 333
Trademarks 335

Index 337

Tables

1. System requirements for the Monitoring Agent for UNIX OS 7
2. Capacity planning for historical data logged by component 115
3. Attributes groups supported by the data model 195
4. Information to gather before contacting IBM Software Support 199
5. Trace log files for troubleshooting agents 202
6. Problems and solutions for installation and configuration 207
7. General problems and solutions for uninstallation 209
8. Agent problems and solutions 211
9. Tivoli Enterprise Portal problems and solutions 216
10. Remote deployment problems and solutions 216
11. Workspace problems and solutions 216
12. Specific situation problems and solutions 218
13. Problems with configuring situations that you solve in the Situation Editor 219
14. Problems with configuration of situations that you solve in the Workspace area 220
15. Problems with configuration of situations that you solve in the Manage Tivoli Enterprise Monitoring Services window 221
16. Take Action commands problems and solutions 221
17. Paging and memory issues for a system administrator to consider 222
18. Time periods and suffixes for summary tables and views. 226
19. Additional columns to report summarization information 226
20. Primary key and warehouse changes for the Monitoring Agent for UNIX OS 228
21. Scripts for affected attribute groups and summary tables for the Monitoring Agent for UNIX OS 229
22. AMS_Pool attribute group (table name: KPX_AMS_POOL) 235
23. CPU_Detail attribute group (table name: KPX_CPU_DETAIL) 235
24. CPU_Summary attribute group (table name: KPX_CPU_SUMMARY) 235
25. Defined_Users attribute group (table name: KPX_DEFINED USERS) 236
26. Devices attribute group (table name: KPX_DEVICES) 237
27. Disks attribute group (table name: KPX_DISKS) 237
28. File_System attribute group (table name: KPX_FILE_SYSTEMS) 237
29. Logical_Partition attribute group (table name: KPX_LOGICAL_PARTITION) 238
30. Network_Adapters_Rates attribute group (table name: KPX_NETWORK_ADAPTERS_RATES) . . . 239
31. Paging_Space attribute group (table name: KPX_PAGING_SPACE) 239
32. Physical_Memory attribute group (table name: KPX_PHYSICAL_MEMORY) 239
33. Processes_Detail attribute group (table name: KPX_PROCESSES_DETAIL). 240
34. Virtual_Memory_Management attribute group (table name: KPX_VIRTUAL_MEMORY_MANAGEMENT). 240
35. WPAR_CPU attribute group (table name: KPX_WPAR_CPU) 240
36. WPAR_FileSystem attribute group (table name: KPX_WPAR_FILESYSTEM) 241
37. WPAR_Information attribute group (table name: KPX_WPAR_INFORMATION) . . . 242
38. WPAR_Network attribute group (table name: KPX_WPAR_NETWORK) 243
39. WPAR_Physical_Memory attribute group (table name: KPX_WPAR_PHYSICAL_MEMORY) 243
40. Overview of Distributed Monitoring migrated situations 245
41. Overview of attribute groups to event classes and slots 248
42. UNIXALLUSR table column heads and the corresponding All Users attributes 275
43. UNIXAMS table column heads and the corresponding AIX AMS attributes 275
44. UNIXCPU table column heads and the corresponding SMP CPU attributes 276
45. UNIXDEVIC table column heads and the corresponding AIX Devices attributes . . . 277
46. UNIXDISK table column heads and the corresponding Disk Information attributes . . 277
47. UNIXDPERF table column heads and the corresponding Disk Performance attributes . 278
48. UNIXDUSERS table column heads and the corresponding AIX Defined Users attributes . 278
49. UNIXFILE table column heads and the corresponding File Information attributes . . 279
50. UNIXGROUP table column heads and the corresponding Group attributes 279
51. UNIXIPADDR table column heads and the corresponding IP Address attributes 279
52. UNIXLPAR table column heads and the corresponding AIX LPAR attributes 280
53. UNIXMACHIN table column heads and the corresponding Machine Information attributes 281
54. UNIXMEM table column heads and the corresponding UNIX Memory attributes . . 281
55. UNIXNET table column heads and the corresponding Network attributes. 283

56. UNIXNSF table column heads and the corresponding NFS and RPC Statistics attributes 284
57. UNIXOS table column heads and the corresponding System attributes 285
58. UNIXPING table column heads and the corresponding Ping attributes 288
59. UNIXPRINTQ table column heads and the corresponding Print Queue attributes . . . 288
60. UNIXPS table column heads and the corresponding Process attributes 289
61. UNIXSOLZON table column heads and the corresponding Solaris Zones attributes . . . 291
62. UNIXTCP table column heads and the corresponding TCP Statistics attributes . . . 291
63. UNIXUSER table column heads and the corresponding User attributes 292
64. UNIXWPARCP table column heads and the corresponding AIX WPAR CPU attributes . . 292
65. UNIXWPARFS table column heads and the corresponding AIX WPAR File System attributes 293
66. UNIXWPARIN table column heads and the corresponding AIX WPAR Information attributes 293
67. UNIXWPARNE table column heads and the corresponding AIX WPAR Network attributes 294
68. UNIXWPARPM table column heads and the corresponding AIX WPAR Physical Memory attributes 294
69. Mechanisms used to gather FILEINFO attributes 297
70. Mechanisms used to gather UNIXALLUSR attributes 297
71. Mechanisms used to gather UNIXAMS attributes 298
72. Mechanisms used to gather UNIXCPU attributes 298
73. Mechanisms used to gather UNIXDEVIC attributes 300
74. Mechanisms used to gather UNIXDISK attributes 300
75. Mechanisms used to gather UNIXDPERF attributes 301
76. Mechanisms used to gather UNIXDUSERS attributes 302
77. Mechanisms used to gather UNIXFILCMP attributes 302
78. Mechanisms used to gather UNIXFILPAT attributes 302
79. Mechanisms used to gather UNIXGROUP attributes 302
80. Mechanisms used to gather UNIXIPADDR attributes 302
81. Mechanisms used to gather UNIXLPAR attributes 303
82. Mechanisms used to gather UNIXMACHIN attributes 304
83. Mechanisms used to gather UNIXMEM attributes 305
84. Mechanisms used to gather UNIXNET attributes 307
85. Mechanisms used to gather UNIXNFS attributes 309
86. Mechanisms used to gather UNIXOS / SP2OS attributes 311
87. Mechanisms used to gather UNIXPING attributes 315
88. Mechanisms used to gather UNXPRINTQ attributes 315
89. Mechanisms used to gather UNIXPS attributes 315
90. Mechanisms used to gather UNIXSOLZON attributes 317
91. Mechanisms used to gather UNIXTCP attributes 318
92. Mechanisms used to gather UNIXUSER attributes 318
93. Mechanisms used to gather UNIXWPARCP attributes 318
94. Mechanisms used to gather UNIXWPARFS attributes 319
95. Mechanisms used to gather UNIXWPARIN attributes 319
96. Mechanisms used to gather UNIXWPARNE attributes 321
97. Mechanisms used to gather UNIXWPARM attributes 322

Chapter 1. Overview of the Monitoring Agent for UNIX OS

The Monitoring Agent for UNIX OS provides you with the capability to monitor and perform basic actions on AIX®, Solaris, and HP-UX operating systems. This chapter provides a description of the features, components, and interface options for the Monitoring Agent for UNIX OS.

IBM Tivoli Monitoring overview

IBM Tivoli Monitoring is the base software for the Monitoring Agent for UNIX OS. IBM Tivoli Monitoring provides a way to monitor the availability and performance of all the systems in your enterprise from one or several designated workstations. It also provides useful historical data that you can use to track trends and to troubleshoot system problems.

You can use IBM Tivoli Monitoring to do the following:

- Monitor for alerts on the systems that you are managing by using predefined situations or custom situations.
- Establish your own performance thresholds.
- Trace the causes leading to an alert.
- Gather comprehensive data about system conditions.
- Use policies to perform actions, schedule work, and automate manual tasks.

The Tivoli Enterprise Portal is the interface for IBM Tivoli Monitoring products. By providing a consolidated view of your environment, the Tivoli Enterprise Portal permits you to monitor and resolve performance issues throughout the enterprise.

See the IBM Tivoli Monitoring publications listed in "IBM Tivoli Monitoring library" on page 327 for complete information about IBM Tivoli Monitoring and the Tivoli Enterprise Portal.

Features of the Monitoring Agent for UNIX OS

The Monitoring Agent for UNIX OS offers a central point of management for your UNIX server environment. This monitoring agent provides a way to monitor the availability and performance of all the systems in your enterprise from one or several designated workstations. This monitoring agent also provides useful historical data that you can use to track trends and to troubleshoot system problems. Information is standardized across all systems (AIX, HP-UX, and Solaris).

The Monitoring Agent for UNIX OS lets you easily collect and analyze server-specific information, such as the following:

- Operating system and CPU performance
- UNIX disk information and performance analysis
- Process status analysis
- Network performance

The Monitoring Agent for UNIX OS provides the following benefits:

- Simplifies application and system management by managing applications, platforms, and resources across your system.
- Increases profits by providing you with real-time access to reliable, up-to-the-minute data that allows you to make faster, better informed operating decisions.
- Scales and ports to a wide variety of UNIX platforms.
- Enhances system performance because you can integrate, monitor, and manage your environment, networks, console, and mission-critical applications. For example, the Monitoring Agent for UNIX OS can alert you when a condition in your environment meet or exceed the thresholds you set. These alerts notify your system administrator to limit and control system traffic. You can view data gathered by the Monitoring Agent for UNIX OS in reports and charts that inform you of the status of your managed UNIX systems.
- Enhances efficiency by monitoring diverse platforms and networks. Depending on the configuration of this monitoring agent, you can collect and monitor data across platforms. The Monitoring Agent for UNIX OS gathers and filters status information at the managed system rather than at the Hub, eliminating unnecessary data transmission and sending only data that is relevant to changes in status conditions.

New in this release

For version 6.2.3 Fixpack 1 of the Monitoring Agent for UNIX OS, the following enhancements have been made:

- A variety of metrics have been ported from the AIX Premium agent to the Monitoring Agent for UNIX OS. New attribute groups include AIX AMS, AIX Defined Users, AIX Devices, AIX LPAR, AIX WPAR CPU, AIX WPAR File System, AIX WPAR Information, AIX WPAR Network, and AIX WPAR Physical Memory. Note that the data collection for the AIX Defined Users attribute group is by default disabled for performance reasons. You must set KUX_DEFINED_USERS=True in the `ux.ini` file to enable it. To view the variety of metrics from the AIX Premium agent, see Appendix B, "AIX Premium agent attributes," on page 235. The minimum AIX requirements to collect the new metrics are
  - AIX53S = AIX 5.3 TL12
  - AIX61F = AIX 6.1 TL5

  Note that the new metrics are not collected unless your environment meets these minimum AIX requirements.
- A new attribute, Volume Group Name (AIX), has been added to the Disk attribute group.
- New attributes, Average Service Queue Size (AIX), Parent (AIX), Service Queue Full per Sec (AIX), Transfers KB per Sec (AIX), and Type (AIX), have been added to the Disk Performance attribute group.
- A new attribute, Bandwidth Utilization Percent (AIX), has been added to the Network attribute group.
- New attributes, Page Space Used (AIX), Resident Data Size (AIX), Resident Text Size (AIX), WLM Name (AIX), and WPAR Name (AIX), have been added to the Process attribute group.
- New attributes, Context Switches per Sec (AIX), Logical Context Switches (AIX), and Physical Consumption (AIX), have been added to the SMP CPU attribute group.

- New attributes, Number of CPUs (AIX), Physical Consumption (AIX), Stolen Busy Cycles Pct (AIX), Stolen Idle Cycles Pct (AIX), System Software Version (AIX), Time Spent in Hypervisor Pct (AIX), have been added to the System attribute group.
- New attributes, Computational Memory (AIX), Decay Rate (AIX), Net Memory Available (MB), Net Memory Available (Percent), Net Memory Used (MB), Net Memory Used (Percent), Non Computational Memory (AIX), Pages Read per Sec (AIX), Pages Written per Sec (AIX), Paging Space Free Pct (AIX), Paging Space Read per Sec (AIX), Paging Space Used Pct (AIX), Paging Space Written per Sec (AIX), Repaging Rate (AIX), and ZFS ARC Size (MB) have been added to the UNIX Memory attribute group.
- The following workspaces:
  - AIX Devices Status
  - AIX LPAR Information
  - AIX WPAR Summary
  - AIX WPAR CPU Details
  - AIX WPAR Details
  - AIX WPAR Memory Details
  - AIX WPAR Network and File System
- A Defined Users view, in the Users workspace, to display information about the defined users, including logon user names, defined roles, whether the user accounts have been locked, expiration dates, and the number of incorrect logon attempts before the account is locked.
- The following situations:
  - UNIX_AIX_Avg_ReqInWaitQ_MS_Info
  - UNIX_AIX_Avg_Transfer_MS_Info
  - UNIX_AIX_CPU_CtxSwitch_Hi_Info
  - UNIX_AIX_Device_Stopped_Warning
  - UNIX_AIX_Memory_RePg_Hi_Info
  - UNIX_AIX_NetBandwidth_High_Info
  - UNIX_AIX_Process_ResDat_Hi_Info
  - UNIX_AIX_Process_ResTxt_Hi_Info
  - UNIX_AIX_ServQ_Full_PerSec_Info
  - UNIX_AIX_System_HypPct_Hi_Info
  - UNIX_AIX_System_NProcs_Hi_Info
  - UNIX_AIX_User_Acct_Locked_Info
  - UNIX_AIX_User_Login_Retry_Info
  - UNIX_LPARBusy_pct_Warning
  - UNIX_LPARPhyBusy_pct_Warning
  - UNIX_LPARvcs_Info
  - UNIX_LPARfreepool_Warning
  - UNIX_LPARPhanIntrs_Info
  - UNIX_LPARentused_Info
  - UNIX_LPAR_MaxCPUCapUsed_Info
  - UNIX_LPAR_Moved_Info
  - UNIX_WPAR_Admin_Op_Info
  - UNIX_WPAR_Broken_Warning

- UNIX_WPAR_CPU_Usage_Warning
- UNIX_WPAR_Mem_Usage_Warning
- UNIX_WPAR_Min_CPU_Limit_Info
- UNIX_WPAR_Min_Mem_Limit_Info
- UNIX_WPAR_RC_Inactive_Info
- UNIX_WPAR_Unlim_CPU_Shares_Info
- UNIX_WPAR_Unlim_Mem_Shares_Info

Note: These predefined situations are based on best practices. While they might not prove perfectly suited to every monitoring environment, they offer a useful starting point for many users.

- By default, when running on AIX 5.3 or later, the kuxagent main spawns a new process. This process, aixdp_daemon, gathers all the metrics available from the kpx data provider and passes them back to kuxagent. By setting the environment variable KUX_AIXDP=false in the `ux.ini` file, an administrator has the option to prevent the aixdp_daemon process from starting. Specify this option when new metrics are not relevant or if performance issues arise. The variable's default value is true (that is, data collection is enabled by default). You can use the ITM V623 remote agent configuration feature to change the value.

For version 6.2.3 of the Monitoring Agent for UNIX OS, the following enhancements have been made:

- Support for self-describing agents. See the *IBM® Tivoli® Monitoring Installation and Setup Guide* for more information.
- Enhanced reporting capabilities, including a redesigned installer for OS Agent reports and new reports for Tivoli Common Reporting. See Chapter 8, “Tivoli Common Reporting for the monitoring agent,” on page 139 for additional informatiton about reporting capabilities.
- A new Tivoli Monitoring capability allows you to perform prerequisite checking for agents before performing an installation. The two mechanisms available are a manually executed stand-alone prerequisite scanner, or a remote prerequisite scanner facility that extends the capabilities of IBM Tivoli Monitoring's remote deployment component. See the *IBM Tivoli Monitoring: Installation and Setup Guide* and the *IBM Tivoli Monitoring: Command Reference* for more information.
- Ability to remotely deploy application agents that support workload partitioning (WPAR) environments. To use this capability, install the Monitoring Agent for UNIX OS on all the WPARs that will be targets of remote deployment operations. Then run the **tacmd addsystem** command, specifying as the node the WPAR managed system.
- Additional filtering capabilities (exploiting regular expressions) on names of processes (full command line). This enhancement is useful for those names that are truncated because of the 768 character limit. For more information, see “Filtering capabilities on the names of processes” on page 12.
- Capability to be alerted about a space availability shortage in the environment, including disks that are assigned to GPFS (general parallel file systems) file systems. This capability includes the File System Type attribute.
- File System Type attribute in the Disk attribute group. Note that any query, applied to an agent prior to V6.2.3, containing this attribute displays a blank in the File System Type column.
- The following attributes in the Solaris Zones attribute group:
  - Capped CPU
  - Capped Memory

  - Dedicated CPU
- TCP Statistics attribute group
- Support of monitoring file systems of the NFS (network file system) type to monitor file systems on systems where the OS agent is not installed, but that are NFS mounted from systems where the OS agent is present.
  - The following environment variable is introduced to enable or disable NFS monitoring at the agent side:

    `KBB_SHOW_NFS=true|false`

    The default value is false.
  - The following environment variable is introduced to enable a timeout:

    `KBB_NFS_TIMEOUT`

    The default value of 2 seconds. The allowed range is 1-30 seconds.
- A new attribute, File System Status, has been added to the Disk attribute group.
- A new, out-of-the-box query, Disk Utilization (623), contains the new File System Status attribute, as well as the recently introduced File System Type attribute.
- The Disk Utilization view in the Disk Usage workspace now contains two new columns:
  - File System Status
  - File System Type
- The following situations:
  - UNIX_BP_AvgCpuBusyPct5min_Criti
  - UNIX_BP_CpuBusyPct_Critical
  - UNIX_BP_LoadAvg5min_Critical
  - UNIX_BP_NetInOutErrPct_Critical
  - UNIX_BP_NumberZombies_Warning
  - UNIX_BP_PagingRate_Critical
  - UNIX_BP_ProcHighCpu_Critical
  - UNIX_BP_ProcMissing_Critical
  - UNIX_BP_SpaceUsedPct_Critical
  - UNIX_BP_SpaceUsedPctCustom_Crit
  - UNIX_BP_SwapSpaceUsedPct_Critic
  - UNIX_BP_SysWaitIOPct_Warning

  Note: These predefined situations are based on best practices. While they might not prove perfectly suited to every monitoring environment, they offer a useful starting point for many users.

Monitoring Agent for UNIX OS components

After you install the Monitoring Agent for UNIX OS (product code "kux" or "ux") as directed in the *IBM Tivoli Monitoring Installation and Setup Guide*, you have an environment with a client, server, and monitoring agent implementation for IBM Tivoli Monitoring that contains the following components:

- Tivoli Enterprise Portal client with a Java-based user interface for viewing and monitoring your enterprise.
- Tivoli Enterprise Portal Server that is placed between the client and the Tivoli Enterprise Monitoring Server and enables retrieval, manipulation, and analysis of data from the monitoring agents.

- Tivoli Enterprise Monitoring Server, which acts as a collection and control point for alerts received from the monitoring agents, and collects their performance and availability data.
- Monitoring Agent for UNIX OS, which collects and distributes data to a Tivoli Enterprise Monitoring Server. This component also embeds the Agent Management Services function.
- Operating system agents and application agents installed on the systems or subsystems you want to monitor. These agents collect and distribute data to the Tivoli Enterprise Monitoring Server.
- Tivoli Data Warehouse for storing historical data collected from agents in your environment. The data warehouse is located on a DB2®, Oracle, or Microsoft SQL database. To collect information to store in this database, you must install the Warehouse Proxy agent. To perform aggregation and pruning functions on the data, install the Warehouse Summarization and Pruning agent.
- Tivoli Enterprise Console event synchronization component for synchronizing the status of situation events that are forwarded to the event server. When the status of an event is updated because of IBM 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. For more information, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

User interface options

Installation of the base software and other integrated applications provides the following interfaces that you can use to work with your resources and data:

Tivoli Enterprise Portal browser client interface
The browser interface is automatically installed with Tivoli Enterprise Portal. To start Tivoli Enterprise Portal in your Internet browser, enter the URL for a specific Tivoli Enterprise Portal browser client installed on your Web server.

Tivoli Enterprise Portal desktop client interface
The desktop interface is a Java-based graphical user interface (GUI) on a Windows workstation.

IBM Tivoli Enterprise Console
Event management application

Manage Tivoli Enterprise Monitoring Services window
The window for the Manage Tivoli Enterprise Monitoring Services utility is used for configuring the agent and starting Tivoli services not already designated to start automatically.

Chapter 2. Requirements for the monitoring agent

This chapter contains information about the requirements for the Monitoring Agent for UNIX OS.

In addition to the requirements described in the *IBM Tivoli Monitoring Installation and Setup Guide*, the Monitoring Agent for UNIX OS has the requirements listed in Table 1.

Table 1. System requirements for the Monitoring Agent for UNIX OS

Operating system	**UNIX**
Operating system versions	• AIX V5.2, 5.3 (32-bit and 64-bit) • AIX V6.1 (64-bit) • HP-UX 11i v1 PA-RISC (32-bit)[1] • HP-UX 11i v2 PA-RISC (64-bit) • HP-UX 11i v3 PA-RISC (64-bit) • HP-UX 11i v2 Integrity (64-bit) • HP-UX 11i v3 Integrity (64-bit) • Solaris V8 on SPARC 32-bit (requires Solaris patches 108434-17, 111721-04, and 109147-07) • Solaris V8 on SPARC 64-bit (requires Solaris patches 108435-17, 111721-04, and 108434-17) • Solaris V9 on SPARC 32-bit (requires Solaris patches 111711-11 and 111722-04) • Solaris V9 on SPARC 64-bit (requires Solaris patches 111712-11, 111722-04, and 111711-11) • Solaris V10 on SPARC (32-bit and 64-bit)[2] • Solaris V10 on x86-64 (64-bit)[2]
Memory	• 55 MB RAM for the Monitoring Agent for UNIX OS
Disk space	The Monitoring Agent for UNIX OS needs 250 MB of disk space in the file system where it is to be installed through the local install method. It needs 184 MB of disk space in the `/tmp` file system and 294 MB of disk space in the file system where the agent is to be installed through the **tacmd createNode** command. It needs 380MB of disk space when it is updated using the command **tacmd updateAgent**. Historical data space varies, depending on the tables collected. Refer to general installation guidelines for disk space requirements in the *IBM Tivoli Monitoring Installation and Setup Guide* and the section in this guide, “Disk capacity planning for historical data” on page 115. Disk space requirements can be as high as 1 GB for log files, remote deploys and historical data.

Table 1. System requirements for the Monitoring Agent for UNIX OS (continued)

Operating system	**UNIX**
Other requirements	• IBM Tivoli Monitoring v6.2.2 agents require at least a v6.2.2 hub monitoring server and portal server. IBM Tivoli Monitoring v6.2.1 hub monitoring servers and portal servers do not support v6.2.2 monitoring agents. IBM Tivoli Monitoring v6.2.1 monitoring agents work with both v6.2.1 and v6.2.2 environments. • A POSIX-compliant threads package must be installed on the monitored machine. • Ethernet or token ring LAN capability. • Native X-term monitor for UNIX or Hummingbird Exceed X-windows emulators for PCs only. • For AIX: A compatible version of libperfstat. For example, AIX 5.2 requires libperfstat V5.2. Upgrade to the latest version of libperfstat for the latest memory fixes. • Most AIX versions require version 8 of the AIX XL C/C++ runtime. For AIX 5.3. you might need version 9. To determine the current level, run the following AIX command: `lslpp -l | grep -i xlc` • The monitoring agent must have the permissions necessary to perform requested actions. For example, if the user ID you used to log onto the system to install the monitoring agent (locally or remotely) does not have the permission to perform a particular action being monitored by the monitoring agent (such as running a particular command), the monitoring agent will be unable to perform the requested action. • Solaris versions require the latest version of SUNWlibC (libC.so.5). • Veritas VxFS (type 32) is a supported file system for the AIX platform.

Notes:

1. As a prerequisite for the Monitoring Agent for UNIX on HP-UX 11i v1, you should have the minimum patch set:
   - PHCO_35743 (PHCO_24400) 'libc cumulative patch'
   - PHCO_31923 (PHCO_24402) 'libc cumulative header file patch'
   - PHCO_32544 (PHCO_23251) 'libc manpage cumulative patch'

   The patches referenced in parenthesis are a minimum level, and the ones on the left of the parenthesis are the latest versions.
2. For Solaris 10, the minimum software group required to run this monitoring agent is the 'End User' group.

Note: For the most current information about the operating systems that are supported, see the following URL: http://publib.boulder.ibm.com/infocenter/prodguid/v1r0/clarity/index.html.

When you get to that site, click on the relevant link in the **Operating system reports** section.

Silent installation: If you are performing a silent installation using a response file, see the *IBM Tivoli Monitoring Installation and Setup Guide*, "Performing a silent installation of IBM Tivoli Monitoring."

Enabling the Monitoring Agent for UNIX OS to run as a nonroot user

The "Post-installation steps for nonroot installations" section of the *IBM Tivoli Monitoring Installation and Setup Guide* describes the post-installation setup process required to enable a nonroot user. Those instructions result in the availability of root authority to the underlying IBM Tivoli Monitoring processes. The instructions in this section, in contrast, remove root authority from the underlying processes.

Securing your IBM Tivoli Monitoring installation

On UNIX operating systems, the product installation process creates the majority of directories and files with **world write** permissions. This configuration creates a security situation that is not acceptable in many enterprises. The **secureMain** utility helps you bring the monitoring environment into compliance with the security standards of your company. Run the **secureMain** utility on all installations, especially those installations that include the UNIX OS Agent, to prevent privilege escalation.

For information about the **secureMain** utility and usage examples, see the "Securing your IBM Tivoli Monitoring installation on Linux or UNIX" appendix in the *IBM Tivoli Monitoring Installation and Setup Guide*.

Setting overall file ownership and permissions for nonroot users

The Monitoring Agent for UNIX OS is capable of running with nonroot user privileges, with some limitations, by changing some agent file permissions and assuring that the desired running user ID has write access to the necessary directories.

The Monitoring Agent for UNIX OS must run with root user privileges to assure correct remote deployment, and collection of some attributes on the Solaris platform. To ensure root privileges, the IBM Tivoli Monitoring installation sets the owner to root and the Set User-ID bit on the primary agent binary, kuxagent, to ensure the agent starts up as the root regardless of which user ID starts the agent.

If you want to start the Monitoring Agent for UNIX OS with permissions of another user ID, use the **chmod** command to turn off the Set User-ID (SUID) bits of the kuxagent binaries to enable running the agent as nonroot. The relevant binary for the Monitoring Agent for UNIX OS in the directory CANDLEHOME/*platform*/ux/bin directory is kuxagent (HPUX - User SUID, Solaris - User SUID, AIX).

Setting kuxagent binary permissions

Changing the permissions requires running systems commands locally on the target system:

```
find CANDLEHOME/* -name kuxagent -exec chmod 755{} \;
```

The bit setting above (755) unsets the SUID bit and ensures that the other bits are set correctly. Note that the bit setting for kuxagent is not persistent. If you ever run secureMain, SetPerm, or install.sh, you need to unset the SUID bit for kuxagent again.

Limitations of starting the agent as a nonroot user

On installation of any other agent by a nonroot user, the permissions on the agent are reset to run the agent with root requirements. You must manually reset the permissions as described above.

Metrics belonging to the WPAR attribute groups:

Note that all of the metrics belonging to the WPAR attribute groups are collected by using the **lswpar** command. However, only the root user can run this command. Therefore, to collect metrics for the WPAR attribute groups, you must be logged into the system as the root user.

Metrics belonging to the Defined Users attributes group:

All of the metrics belonging to the Defined Users attribute group are collected by using the **lsuser -c ALL** command. To collect metrics for the Defined Users attribute group as a nonroot user, you must belong to the security group. If not, the Defined Users view of the Users workspace lists "Not Collected" for each of its fields. In addition, even if the user belongs to the security group, the Roles and Login Retries attributes of the Defined Users group might be incorrectly reported as Not Collected.

Remote Deployment:

Remote deployment might not complete or work at all on certain agents that require root privileges to install the desired application. Install the agents locally or configure the agent manually after installation.

Take Action commands:

The agent cannot start any Take Action commands that require privileged permissions that the user ID does not have.

Solaris:

While running as a nonroot user, the agent cannot access `/proc/pid/status`, and therefore cannot report the following attributes:

- User CPU Time (UNIXPS.USERTIME)
- System CPU Time (UNIXPS.SYSTEMTIM)
- Total CPU Time (UNIXPS.TOTALTIME)
- Thread Count (UNIXPS.THREADCNT)
- Child User CPU Time (UNIXPS.CHILDUTIME)
- Child System CPU Time (UNIXPS.CHILDSTIME)
- Total Child CPU Time (UNIXPS.CHILDTIME)
- Wait CPU Time (UNIXPS.WAITCPUTIM)
- Terminal (UNIXPS.USERTTY)

Agent Management Services:

Data reported in the Agent Active Runtime Status attribute group, for example the PID, the command line, the CPU, and the memory, might also be affected when the nonroot user is monitoring agents running as a different nonroot user. The watchdog cannot stop or start any agent that it does not have privileges to stop or

start. See “Using Agent Management Services” for a complete description of the Agent Watchdog and Agent Management Services Watchdog.

Setting up the Monitoring Agent for UNIX OS in a cluster environment

The *IBM Tivoli Monitoring Installation and Setup Guide* contains an overview of clustering. The information provided here is specifically for installing and setting up the Monitoring Agent for UNIX OS in an HACMP™ clustered environment.

The Monitoring Agent for UNIX OS is set up and works as it does in a non-clustered environment. There is a unique cluster configuration prompt for the Monitoring Agent for UNIX OS:

```
Are you installing this product into a clustered environment (Default is: NO):
```

You should accept the default (NO).

Note: The NO or YES response is case-sensitive.

The shared disks attributes are displayed in all cluster nodes workspace views that use the disk-related queries. However, the disk metrics are zero for the nodes that are not controlling the shared disk because they do not have access to the shared disk. The highly available IP addresses are not displayed in the network table since the highly available IP address is bound to an alias interface. Alias interface IP addresses are reported in the IP Address attribute group. The relevant statistics for those IP addresses are reported in the base Interface name in the Network attribute group since all aliases share one set of statistics for a particular network interface. The base interface can usually be found by removing the colon and number at the end of the aliased interface (for example, Alias = en0:3 Base: en0).

Using Agent Management Services

There are two watchdog monitors that run as part of the Monitoring Agent for UNIX. One monitor runs as part of the OS Monitoring Agent process, which is referred to as the *Agent Watchdog*. The other watchdog monitor runs as a separate process named 'kcawd'. The kcawd process is also called the *Agent Management Services Watchdog*. This watchdog monitor watches the OS Agent as long as its Availability Status is showing 'Running' in the Agent's Runtime Status view of the Agent Management Services workspace. No setup or configuration is required.

The Agent Watchdog monitors agent processes other than the OS Agent. By using the communication facility of the OS Agent, the monitor can respond to Tivoli Enterprise Portal Desktop queries and Take Action commands that are performed against these other agent processes. This data is displayed in the Agent Management Services workspace. In the Tivoli Enterprise Portal Desktop, the Agent Management Services workspace lists the agents that can be monitored by this watchdog that is running as part of the OS Agent. These agents are non-OS agents, so the Monitoring Agent for UNIX is not listed in the workspace, except for in the Agents' Management Definitions view. One of the agents listed in the workspace is the Agent Management Services Watchdog. Its purpose is to monitor the OS Agent's availability.

The Agent Management Services Watchdog monitor is responsible for watching just the OS Monitoring Agent and restarting it if it goes down. It is enabled by default and does not need to be configured. It is started automatically when the Monitoring Agent for UNIX is started. This watchdog does not have a communication facility, so it cannot report information to the Tivoli Enterprise

Portal or respond to Take Action commands. It is not an agent in itself, but a separate process that always monitors the OS Monitoring Agent.

You can temporarily disable the Agent Management Services Watchdog by using the *InstallDir*/bin/itmcmd execute ux disarmWatchdog.sh command. This command disables the Watchdog process for the OS Monitoring Agent and all Agent Management Services managed agents. If there is local administrative work to be performed, and you do not want the auto-restart of the agents to interfere with it, run the *InstallDir*/bin/itmcmd execute ux disarmWatchdog.sh command before proceeding. When the work is complete, recycle the OS Monitoring Agent to reenable Agent Management Services, or use the *InstallDir*/bin/itmcmd execute ux rearmWatchdog.sh command.

If you use the itmcmd interface to stop or start an Agent Management Services managed agent, its watchdog will be disabled if stopping the agent and enabled if starting the agent.

Filtering capabilities on the names of processes

You can now distinguish process names that are longer than 768 characters, so that situations can be defined on the relevant part of the name. You can also use this enhancement for filtering processes of any length.

To improve filtering on the processes, a Process Filter attribute has been added to the UNIX Process attribute group. Its content, a regular expression, is sent to the agent as a filter object and is intended to only act on the Process Command (Unicode) attribute. For example, the agent uses the value provided in the Process Filter attribute to match the process name, and then fills the Process Command (Unicode) attribute.

In a Tivoli Enterprise Portal workspace view, you see only the processes whose names match the specified regular expression. The Process Command (Unicode) column is filled with the matching patterns separated by blanks, as defined in the regular expression. The Process Filter column is filled with the regular expression that matches it. Situations can be defined mixing the Process Command (Unicode) column and other conditions (for example, CPU usage).

To use this enhancement, create queries and situations on the UNIX Process attribute group containing the Process Filter attribute and define a regular expression in it. More rows and more regular expressions are allowed. Use the query in a workspace view or distribute the situation to the target managed systems.

There are a few predefined regular expressions for the Process Filter attribute when you use it in the query or situation editor:

- Java processes (.*java.*)
- IBM_Java_processes_entry_method_only_(.*java.*(com.ibm.*))
- System Admin installed processes_(/usr.*)

Chapter 3. Workspaces reference

This chapter contains an overview of workspaces, references for detailed information about workspaces, and descriptions of the predefined workspaces included in this monitoring agent.

About workspaces

A workspace is the working area of the Tivoli Enterprise Portal application window. At the left of the workspace is a Navigator that you use to select the workspace you want to see.

As you select items in the Navigator, the workspace presents views pertinent to your selection. Each workspace has at least one view. Some views have links to workspaces. Every workspace has a set of properties associated with it.

This monitoring agent provides predefined workspaces. You cannot modify the predefined workspaces, but you can create new workspaces by editing them and saving the changes with a different name.

More information about workspaces

For more information about creating, customizing, and working with workspaces, see the *IBM Tivoli Monitoring User's Guide*.

For a list of the predefined workspaces for this monitoring agent and a description of each workspace, refer to the Predefined workspaces section below and the information in that section for each individual workspace.

Predefined workspaces

The Monitoring Agent for UNIX OS provides the following predefined workspaces, which are organized by Navigator item:

The following list shows how the IBM Tivoli Monitoring: UNIX OS Agent workspaces are organized.

- "UNIX workspace" on page 31
  - "UNIX Detail workspace" on page 31
  - "Enterprise UNIX System Summary workspace" on page 19
- "Agent Management Services workspace" on page 15
  - "Agents' Management Log workspace" on page 15
- "Disk Usage workspace" on page 18
  - "Disk Usage Details workspace" on page 19
  - "Disk Utilization for Mount Point workspace" on page 19
- "File Information workspace" on page 20
  - "All Files workspace" on page 18
- "Network workspace" on page 26
- "NFS Activity workspace" on page 27
- "Process workspace" on page 28

  - "All Processes workspace" on page 18
  - "Top CPU-Memory %-VSize Details workspace" on page 31
  - "Solaris Zone Processes workspace" on page 29
- "RPC Performance workspace" on page 29
- "System Information workspace" on page 30
  - "Solaris System CPU Workload workspace" on page 29
  - "System Details workspace" on page 30
  - "Solaris Zones workspace" on page 29
  - "AIX Devices Status workspace" on page 16
  - "AIX LPAR Information workspace" on page 16
  - "AIX WPAR Summary workspace" on page 16
    - "AIX WPAR CPU Details workspace" on page 17
    - "AIX WPAR Details workspace" on page 17
    - "AIX WPAR Memory Details workspace" on page 17
    - "AIX WPAR Network and File System workspace" on page 18
- "Users workspace" on page 32

This agent also includes the following historical workspaces:

- "Historical Summarized Availability workspace" on page 20
- "Historical Summarized Availability Daily workspace" on page 20
- "Historical Summarized Availability Hourly workspace" on page 20
- "Historical Summarized Availability Weekly workspace" on page 21
- "Historical Summarized Capacity workspace" on page 21
- "Historical Summarized Capacity Daily workspace" on page 22
- "Historical Summarized Capacity Hourly workspace" on page 22
- "Historical Summarized Capacity Weekly workspace" on page 23
- "Historical Summarized Performance workspace" on page 24
- "Historical Summarized Performance Daily workspace" on page 24
- "Historical Summarized Performance Hourly workspace" on page 25
- "Historical Summarized Performance Weekly workspace" on page 26

Some predefined workspaces are not available from the Navigator tree item, but are accessed by selecting the link indicator next to a row of data in a view. Left-clicking a link indicator selects the default workspace associated with that link. Right-clicking a link indicator displays all linked workspaces that can be selected. Examples of the workspaces for this monitoring agent include the following:

- Application for Process - which includes the Application for Process ID table view and the Take Action view.
- Child Processes - which includes the Child Processes for Process ID table view and Take Action view.
- Command for Process - which includes the Command for Process ID table view and Take Action view.
- Processes for Group Leader - which includes the Processes for Group Leader ID table view and Take Action view.
- Process Resource - which includes the Resource Usage for Process ID table view and Take Action view. This workspace has a superseded version that displays

queries with signed 32-bit maximum value (2,147,483,647) and a version with the same name (minus 'superseded') with queries that support values up to signed 64-bit max (9,223,372,036,854,775,807).

Agent Management Services workspace

The Agent Management Services workspace contains views of data collected by the Agent Management Services component of the Monitoring Agent for UNIX. This workspace includes an Agents' Management Status view, an Agents' Runtime Status view, an Agents' Alerts view, and an Agents' Management Definitions view.

Agents' Management Log workspace

The Agents' Management Log workspace contains a list of monitoring agent log entries filtered on the Agent Management Services component. Use this workspace to see the operations being executed by Agent Management Services and to keep an audit trail of these operations.

Log messages generated by the physical watchdog are displayed in the Agents' Management Log workspace view. By using these log messages, you can track OS Agent restarts and availability. The limitations of this function are that the physical watchdog must be running.

Alerts that are seen in the Alerts view in the default workspace are cached for 24 hours, by default. The time can be overridden by changing the environment variable KCA_CACHE_LIMIT found in the `ux.ini` file. The variable is specified in hours. This functionality is not available to previous versions of the agents.

The workspace includes the following operation messages:

- Agent added to system - CAP file found.
- Agent CAP file initialization completed.
- Agent daily restart count reset.
- Agent exceeded policy defined CPU threshold.
- Agent exceeded policy defined memory threshold.
- Agent exceeded restart tries.
- Agent initial start.
- Agent Management Services watchdog not reliable.
- Agent manual start failed.
- Agent manual stop failed.
- Agent not configured.
- Agent not found.
- Agent now managed.
- Agent now unmanaged.
- Agent recycle command received.
- Agent removed from system - CAP file removed.
- Agent restart disabled - disarm mode active
- Agent restart failed.
- Agent start command received.
- Agent started successfully.
- Agent stop command received.
- Agent stopped abnormally.

- Agent stopped successfully.
- Disarm completed successfully.
- Rearm completed successfully.

This workspace includes an Agents' Management Log view.

AIX Devices Status workspace

The AIX Devices Status workspace, accessed by right-clicking on the System Information navigator item, contains views of data collected by the AIX Devices component of the Monitoring Agent for UNIX. This workspace includes a Device Status table view, indicating the name, parent name, state, class, and type of the AIX devices.

AIX LPAR Information workspace

The AIX LPAR Information workspace, accessed by right-clicking on the System Information navigator item, contains views of data collected by the AIX LPAR component of the Monitoring Agent for UNIX. This workspace contains the following views:

LPAR CPU Utilization (pie chart)
This chart presents a graphic view of the LPAR CPU utilization based on the percent used by CPU type.

CPU Entitlement (bar chart)
This chart displays the number of physical CPUs and the number of entitlement units assigned to this logical partition.

LPAR CPUs (bar chart)
This chart displays the number of logical CPUs, virtual CPUs, physical CPUs in the shared pool, and available CPUs from the shared pool.

LPAR Attributes (table view)
This table provides a variety of details for the specific logical partition.

LPAR Utilization (table view)
This table offers a comprehensive list of details related to logical partition usage.

Active Memory Sharing (AMS) Pool (table view)
This table offers a comprehensive list of details related to the Active Memory Sharing (AMS) pool.

AIX WPAR Summary workspace

The AIX WPAR workspaces are comprised of a summary workspace and four detail workspaces, each accessible through a link in the WPAR Status and Configuration view on the AIX WPAR Summary workspace. This workspace, accessed by right-clicking on the System Information navigator item, provides a summary of CPU utilization, memory utilization, and current status of all of the WPARs. This workspace contains the following views:

LPAR CPU Utilization by WPARs (bar chart)
Shows the percentage of entitlement for the LPAR consumed by each WPAR.

LPAR Memory Utilization by WPARs (bar chart)
Shows the percentage of memory for the LPAR used by each WPAR.

WPAR Status and Configuration (table view)
Provides a summarized status for all of the WPARs.

AIX WPAR CPU Details workspace

The AIX WPAR CPU Details workspace provides detailed information on CPU consumption and CPU resource limits for a WPAR. This workspace contains the following views:

CPU Usage (bar chart)
Shows the CPU consumption and consumption limit for the WPAR.

CPU Usage Modes (pie chart)
Shows the percentage CPU utilization in user and system modes for the WPAR.

CPU Limit (bar chart)
Shows the CPU consumption limit for the WPAR and the entitlement for the LPAR.

CPU Usage Details (table view)
Shows detailed CPU usage and limits for the WPAR.

Workspace Navigation (table view)
Lists all of the WPARs in the partition with links beside them to navigate to the desired workspace.

AIX WPAR Details workspace

The AIX WPAR Details workspace provides detailed views on the status, configuration, and resource limits of a WPAR. This workspace contains the following views:

Resource Controls (table view)
Shows various resource controls for the WPAR.

CPU Resource Limits (table view)
Lists the CPU resource limits for the WPAR.

Memory Resource Limits (table view)
Lists the memory resource limits for the WPAR.

Configuration (table view)
Shows various configuration parameters for the WPAR.

Administration Status (table view)
Shows a table providing administrative information and status of the WPAR.

AIX WPAR Memory Details workspace

The AIX WPAR Memory Details workspace provides detailed information about memory consumption and memory resource limits for a WPAR. This workspace contains the following views:

Used and Free Memory (bar chart)
Shows chart with the amount of free and used memory in the WPAR.

Memory Utilization (pie chart)
Shows the used and free memory percentages for the WPAR as a chart.

Memory Usage and Availability (bar chart)
Shows memory usage for the WPAR and memory size against the total memory available in the partition.

Memory Usage Details (table view)
Shows detailed memory usage and limits for the WPAR.

Workspace Navigation (table view)
Lists all of the WPARs in the partition with links beside them to navigate to the desired workspace.

AIX WPAR Network and File System workspace

The AIX WPAR Network and File System workspace displays networking and file system information for a WPAR. This workspace contains the following views:

Network Summary (table view)
Shows a table with network interface-related metrics for the WPAR.

File System Summary (table view)
Shows the file system information for the WPAR.

Workspace Navigation (table view)
Lists all of the WPARs in the partition with links beside them to navigate to the desired workspace.

All Files workspace

The All Files workspace can be accessed by right-clicking on the File Information navigator item. The All Files workspace includes the File Size - Top Ten view and the All Files view.

All Processes workspace

The All Processes workspace includes the Top CPU Time, Memory %, Processes, and Virtual Size views, accessed by right-clicking the Process navigator item.

Disk Usage workspace

The Disk Usage workspace is comprised of four views and provides an at-a-glance snapshot of your disk condition. This workspace has a superseded version that displays queries with signed 32-bit maximum value (2,147,483,647) and a version with the same name (minus 'superseded') with queries that support values up to signed 64-bit max (9,223,372,036,854,775,807). This workspace contains the following views:

Space Used % - Top Ten (bar chart)
This chart provides information on the top ten subdirectories using hard drive space.

Inodes Used % - Top Ten (bar chart)
This chart presents a diagrammatic view of the percent of inodes used by the top ten drive space components.

Disk Utilization (table view)
The Disk Utilization report helps you solve disk-related problems quickly by providing information on devices with excessive I/O activity and/or long service times. New data includes average disk queue length and average number of processes waiting for service.

Note: By viewing monitored disk input/output (I/O) data collected by the UNIX agents, you can perform the following tasks:

- Improve the performance of your UNIX system by quickly and accurately pinpointing how your system transfers data from disk to memory

- Increase your efficiency by keeping you constantly informed of how the disk performs

You monitor the performance of disks in your system to make sure they operate smoothly. Disk performance has a major impact on the overall functioning of your UNIX system. If disk performance is poor, it affects jobs that require a large amount of disk I/O, virtual memory functioning, and the time required to run a program.

Use this report to identify disk performance problems caused by slow rates of data transfer from disk to memory, or high disk usage.

Disk Busy % (bar chart)

This graph shows monitored disk input/output (I/O) data collected by the UNIX OS Agents. With this information you can perform the following tasks:

- Improve the performance of your UNIX system by quickly and accurately pinpointing how your system transfers data from disk to memory
- Increase your efficiency by keeping you constantly informed of how the disk performs

Disk Usage Details workspace

The Disk Usage Details workspace includes the Disk Utilization, Disk Performance, and Disk Inode views. This workspace can be accessed by right-clicking the Disk Usage navigator item. This workspace has a superseded version that displays queries with signed 32-bit maximum value (2,147,483,647) and a version with the same name (minus 'superseded') with queries that support values up to signed 64-bit max (9,223,372,036,854,775,807).

Disk Utilization for Mount Point workspace

The Disk Utilization for Mount Point workspace includes the Disk Utilization for Mount Point table view, pie chart view, and Take Action view. This workspace can be accessed by right-clicking the Disk Usage navigator item. This workspace has a superseded version that displays queries with signed 32-bit maximum value (2,147,483,647) and a version with the same name (minus 'superseded') with queries that support values up to signed 64-bit max (9,223,372,036,854,775,807).

Enterprise UNIX System Summary workspace

The Enterprise UNIX System Summary workspace requires that the hub and remote Tivoli Enterprise Monitoring Servers are seeded with the IBM Tivoli Monitoring: UNIX OS Agent seed data. This workspace, accessed by right-clicking the UNIX navigator tree item, displays a summary of all UNIX systems on-line with the following views:

UNIX System Summary (table view)

UNIX System Memory Summary (table view)

Memory Usage Summary (bar chart)

Load Average Summary (bar chart)

CPU % Summary (bar chart)

File Information workspace

The File Information workspace is comprised of two views. This workspace contains the following views:

File Size - Top Ten (bar chart)
This graph allows you to quickly see at a glance the files that are consuming the largest amount of disk space.

File Size - Top Ten (table view)
This table allows you to quickly see at a glance the files that are consuming the largest amount of disk space.

The Specific File Information workspace can be accessed by right-clicking the link on either the File Information workspace or the All Files workspace. The Specific File Information workspace includes the File Information and Take Actions views.

Historical Summarized Availability workspace

The Historical Summarized Availability workspace shows the average availability of managed resources over the number of months that you specify in the Time Span dialog. This workspace consists of the following two graphical views:

- Availability by O/S Type (average over months), which shows the percentage of aggregate time that managed resources were available, grouped by operating system
- Availability by Server (average over months), which shows the percentage of aggregate time that each server was available

Historical Summarized Availability Daily workspace

The Historical Summarized Availability Daily workspace shows the availability of the server and the system summary for a server by day. This workspace consists of the following two graphical views:

- Availability (daily), which shows the total percentage of aggregate time that a server was available, summarized by day
- System Summary (daily), which shows system data such as the following, summarized by day:
  - Internet Protocol (IP) address of a monitored system, expressed in dotted decimal format
  - Average sum of the System CPU and User CPU attributes expressed as a percentage
  - Maximum amount of virtual memory, in kilobytes, that the system was using
  - Maximum percentage of virtual memory that was used
  - Average number of processes that were running on a processor
  - Operating system type and version

Historical Summarized Availability Hourly workspace

The Historical Summarized Availability Hourly workspace shows the availability of the server and the system summary for a server by hour. This workspace consists of the following two graphical views:

- Availability (hourly), which shows the total percentage of aggregate time that a server was available, summarized by hour
- System Summary (hourly), which shows system data such as the following, summarized by hour:

- Internet Protocol (IP) address of a monitored system, expressed in dotted decimal format
- Average sum of the System CPU and User CPU attributes expressed as a percentage
- Maximum amount of virtual memory, in kilobytes, that the system was using
- Maximum percentage of virtual memory that was used
- Average number of processes that were running on a processor
- Operating system type and version

Historical Summarized Availability Weekly workspace

The Historical Summarized Availability Weekly workspace shows the availability of the server and the system summary for a server by week. This workspace consists of the following two graphical views:

- Availability (weekly), which shows the total percentage of aggregate time that a server was available, summarized by week
- System Summary (weekly), which shows system data such as the following, summarized by week:
  - Internet Protocol (IP) address of a monitored system, expressed in dotted decimal format
  - Average sum of the System CPU and User CPU attributes expressed as a percentage
  - Maximum amount of virtual memory, in kilobytes, that the system was using
  - Maximum percentage of virtual memory that was used
  - Average number of processes that were running on a processor
  - Operating system type and version

Historical Summarized Capacity workspace

The Historical Summarized Capacity workspace shows the percentage of system resources used for the time span that you specify in the Time Span dialog. This workspace consists of the following four graphical views:

- Network Capacity (average over months), which shows the following ratio regarding the time span that you specify in the Time Span dialog: `Total number of times that a frame transmitted by the network interface collided with another frame to the total number of frames received and transmitted`

  Note: The Frames_Received and Frames_Transmitted attributes really refer to packets, not frames.
- Processor Utilization (average over months), which shows the average percentage of idle CPU time during the sampling period
- Memory Utilization (average over months), which shows the average and maximum percentages of virtual memory used, ordered by the maximum percentage of virtual memory used, during the specified time period
- Disk Utilization (average over months), which shows the maximum percentage of space used (expressed as a percentage of the sum of used and available space) on all the system's logical disks during the specified time period. This value shows the percentage of disk space that is available to non-superusers.

Historical Summarized Capacity Daily workspace

The Historical Summarized Capacity Daily workspace shows the percentage of system resources used daily. This workspace consists of the following four graphical views:

- Network Capacity (daily), which shows the following information for each network instance, summarized by day:
  - Percentage of total times that a frame transmitted by the network interface collided with another frame
  - Number of frames received since the network interface was configured
  - Total number of eight-bit packets transmitted per second by an interface since boot time
- Processor Utilization (daily), which shows the following information, summarized by day:
  - Average percentage of idle CPU time during the sampling period
  - Percentage of processors' CPU time devoted to running UNIX system kernel code
  - Percentage of processors' CPU time devoted to user processes, including time spent executing user program and library functions but not executing system calls
  - Wait I/O CPU time during the sampling period, which indicates how effectively the system or a processor is using disks
- Memory Utilization (daily), which shows the average and maximum percentage of virtual memory used, ordered by the maximum percentage of virtual memory used, summarized by day
- Disk Utilization (daily), which shows space used on all the system's logical disks, including the following details, summarized by day:
  - Name of the physical disk partition where the file system is mounted (the physical location of the disk)
  - Mount point, the path name of the directory to which a file system was mounted and thus the virtual name for the directory
  - Total size of a file system, expressed in KB
  - Maximum amount of disk space used on a file system, expressed in kilobytes
  - Maximum space used on the file system, expressed as a percentage of the sum of used and available space
  - Minimum available space
  - Maximum percentage of inode space that was allocated to files

Historical Summarized Capacity Hourly workspace

The Historical Summarized Capacity Hourly workspace shows the percentage of system resources used, summarized by hour. This workspace consists of the following four graphical views:

- Network Capacity (hourly), which shows the following information for each network instance, summarized by hour:
  - Percentage of total times that a frame transmitted by the network interface collided with another frame
  - Number of frames received since the network interface was configured
  - Total number of eight-bit packets transmitted per second by an interface since boot time

- Processor Utilization (hourly), which shows the following information for each processor in the server, summarized by hour:
  - CPU ID
  - Average percentage of idle CPU time during the sampling period
  - Average percentage of system or per processor CPU time devoted to running UNIX system kernel code
  - Average percentage of system or per processor CPU time devoted to user processes, including time spent executing user program and library functions but not executing system calls
  - Average wait I/O CPU time during the sampling period, which indicates how effectively the system or a processor is using disks
  - Average number of processes that were waiting to be run or were running on a processor
  - Average number of processes that were idle, or in sleep state, or in a stopped state
  - Average number of zombie processes
- Memory Utilization (hourly), which shows the average and maximum percentage of virtual memory used, ordered by the maximum percentage of virtual memory used, summarized by hour
- Disk Utilization (hourly), which shows space used on all the system's logical disks, including the following details, summarized by hour:
  - Name of the physical disk partition where the file system is mounted (the physical location of the disk)
  - Mount point, the path name of the directory to which a file system is mounted and thus the virtual name for the directory
  - Total size of a file system, expressed in KB
  - Maximum amount of disk space used on a file system, expressed in kilobytes
  - Maximum space used on the file system, expressed as a percentage of the sum of used and available space
  - Minimum available space
  - Maximum number and maximum percentage of inode space that was allocated to files
  - Minimum number of inodes that was available on your file system, which can help you take corrective action such as freeing unneeded space or deleting temporary files

Historical Summarized Capacity Weekly workspace

The Historical Summarized Capacity Weekly workspace shows the percentage of system resources used, summarized by week. This workspace consists of the following four graphical views:

- Network Capacity (weekly), which shows the total number of times that a frame transmitted by the network interface collided with another frame, summarized by week
- Processor Utilization (weekly), which shows the average percentage of idle CPU time, summarized by week
- Memory Utilization (weekly), which shows the average and maximum percentage of virtual memory used, ordered by the maximum percentage of virtual memory used, summarized by week
- Disk Utilization (weekly), which shows space used on all the system's logical disks, including the following details, summarized by week:

- Name of the physical disk partition where the file system was mounted (the physical location of the disk)
- Mount point, the path name of the directory to which a file system was mounted and thus the virtual name for the directory
- Total size of a file system, expressed in KB
- Maximum amount of disk space used on a file system, expressed in kilobytes
- Maximum space used on the file system, expressed as a percentage of the sum of used and available space
- Minimum available space
- Maximum percentage of inode space that was allocated to files

Historical Summarized Performance workspace

The Historical Summarized Performance workspace shows the average performance of system resources for the time span that you specify in the Time Span dialog. This workspace consists of the following five graphical views:

- Network Interface Performance (average over months), which shows (calculated per second) the total number of frames received plus frames transmitted for all network interfaces since the network interface was configured, during the time span that you specify in the Time Span dialog
- Load Average (average over months), which shows the average of the average number of processes in the UNIX kernel run queue during a 15-minute interval, during the specified time period
- Processor Performance (average over months), which shows the average percentage of time that the CPU was not processing instructions, during the specified time period
- Memory Performance (average over months), which shows the average rate of page in and page out for the system during the specified time period
- Disk Performance (average over months), which shows the maximum percentage that the disk was busy during the specified time period

Historical Summarized Performance Daily workspace

The Historical Summarized Performance Daily workspace shows the average performance of system resources, summarized by day. This workspace consists of the following five graphical views:

- Network Interface Performance (daily), which shows the following information for each network interface, summarized by day:
  - Total number of frames received
  - Total number of frames transmitted
  - Total input errors
  - Total output errors
  - Total number of times during the sampling period that a frame transmitted by the network interface collided with another frame
- Load Average (daily), which shows the average of the average number of processes in the UNIX kernel run queue during a 15-minute interval and the average percentage of time that the CPU was not processing instructions, summarized by day
- Processor Performance (daily), which shows the following information, summarized by day:
  - Average percentage of idle CPU time during the sampling period

  - Average percentage of processors' CPU time devoted to running UNIX system kernel code
  - Average percentage of processors' CPU time devoted to user processes, including time spent executing user program and library functions but not executing system calls
  - Average wait I/O CPU time during the sampling period, which indicates how effectively the system or a processor is using disks
- Memory Performance (daily), which shows the average rate of page in and page out for the system, summarized by day
- Disk Performance (daily), which shows the following information, summarized by day:
  - Name of the disk
  - Maximum percentage of time a disk was busy transferring data, which can indicate whether a process is I/O bound
  - Average number of disk requests outstanding during the sampling period
  - Average amount of disk time used in milliseconds over the sampling period
  - Average time waiting for disk access, expressed in milliseconds
  - Maximum number of total kilobytes per second written to and read from all local disks during a specified sampling period

Historical Summarized Performance Hourly workspace

The Historical Summarized Performance Hourly workspace shows the average performance of system resources, summarized by hour. This workspace consists of the following five graphical views:

- Network Interface Performance (hourly), which shows the following information for each network interface, summarized by hour:
  - Total number of frames received
  - Total number of frames transmitted
  - Total input errors
  - Total output errors
  - Total number of times during the sampling period that a frame transmitted by the network interface collided with another frame
- Load Average (hourly), which shows the average of the average number of processes in the UNIX kernel run queue during a 15-minute interval and the average percentage of time that the CPU was not processing instructions, summarized by hour
- Processor Performance (hourly), which shows the following information for each processor, summarized by hour:
  - Average percentage of idle CPU time during the sampling period
  - Average percentage of processors' CPU time devoted to running UNIX system kernel code
  - Average percentage of processors' CPU time devoted to user processes, including time spent executing user program and library functions but not executing system calls
  - Average wait I/O CPU time during the sampling period, which indicates how effectively the system or a processor is using disks
  - Average numbers of processes that were waiting to be run or were running on a processor

  - Average numbers of processes that were idle, or in sleep state, or in a stopped state
  - Average number of zombie processes
- Memory Performance (hourly), which shows the average rate of page in and page out rates for the system, summarized by hour
- Disk Performance (hourly), which shows the following information, summarized by hour:
  - Name of the disk
  - Maximum percentage of time a disk was busy transferring data, which can indicate whether a process is I/O bound
  - Average number of disk requests outstanding during the sampling period
  - Average amount of disk time used in milliseconds over the sampling period
  - Average time waiting for disk access, expressed in milliseconds
  - Maximum number of total kilobytes per second written to and read from all local disks during a specified sampling period

Historical Summarized Performance Weekly workspace

The Historical Summarized Performance Weekly workspace shows the average performance of system resources, summarized by week. This workspace consists of the following five graphical views:

- Network Interface Performance (weekly), which shows (calculated per second) the total number of frames received plus frames transmitted for all network interfaces since the network interface was configured, summarized by week
- Load Average (weekly), which shows the average of the average number of processes in the UNIX kernel run queue during a 15-minute interval, summarized by week
- Processor Performance (weekly), which shows the average percentage of time that the CPU was not processing instructions, summarized by week
- Memory Performance (weekly), which shows the average rate of page in and page out for the system, summarized by week
- Disk Performance (weekly), which shows the maximum percentage that the disk was busy, summarized by week

Network workspace

The Network workspace is comprised of four views. This workspace has a superseded version that displays queries with signed 32-bit maximum value (2,147,483,647) and a version with the same name (minus 'superseded') with queries that support values up to signed 64-bit max (9,223,372,036,854,775,807). This workspace contains the following views:

Errors and collisions (bar chart)

The chart displays network input errors, output errors, and collisions allowing you a quick glance at potentially severe network problems.

Packet Transfers (bar chart)

The Packet Transfers chart displays graphically the frames received and frames transmitted on each network card.

The Received Count (Frames) and Transmitted Count (Frames) show the raw frame counts for the interface. Frames and packets are not necessarily the same thing.

System administrators might define the interface so that multiple frames are sent or received in a packet. The network report and the network attributes display frame counts.

Network (table view)
The Network view contains monitored data that helps you see your overall network performance, and supplies you with information about network interfaces on your monitored UNIX systems.

Using this real-time network information, you can perform the following tasks:

- Fine-tune your network by allowing you to pinpoint bottlenecks in network traffic so your system administrator can reroute jobs to avoid them
- Increase system productivity by providing you with the information that lets you determine how network resources can be used more efficiently

This view contains numerical attributes that report the changing interface workload throughout each workday. This information allows your system administrator to see the ebb and flow of network traffic.

The Network view helps you perform the following tasks:

- See how much data your network interfaces transmit and receive
- Identify network interfaces
- Determine whether the interfaces are operational
- View the amount of data traffic that flows through interfaces
- Justify rerouting network data traffic, if necessary
- Measure errors in frame transmission and collisions

IP Addresses (table view)
The view displays IP Address information.

NFS Activity workspace

The NFS Activity workspace is comprised of four views. This workspace contains the following views:

NFS Server Statistics (bar chart)
The NFS Server Statistics chart graphically displays the activity for the following attributes:

- NFS Server Calls
- NFS Server Calls Rejected
- NFS Server Null Calls
- NFS Server lookups

NFS Client Statistics (bar chart)
The NFS Client Statistics chart graphically displays the activity for the following attributes:

- NFS Client Calls
- NFS Client Calls Rejected
- NFS Client Null Calls
- NFS Client lookups

NFS Server Activity (table view)
NFS Server Activity attributes report on NFS calls to the managed system. The agent reports these calls only when the managed system is an NFS server.

NFS Client Activity (table view)
NFS Client Activity attributes report on calls from the managed system to NFS servers.

NFS RPC Statistics group is a single-instance attribute group. You can mix these attributes with those of any other single-instance group.

Process workspace

The Process workspace is comprised of four views. This workspace contains the following views:

Top CPU Time % Processes (bar chart)
This chart displays CPU Percent for the processes utilizing the highest CPU time.

Memory % - Top Ten (bar chart)
This chart displays graphically the Memory Percent attribute for the top ten processes utilizing memory.

Top CPU Time % Processes (table view)
The Top CPU Time Processes report provides detailed information about the top CPU processes currently running. You can view this information in the detail report columns.

The Top CPU Time Processes report help you perform the following tasks:

- Solve process problems quickly by giving you information you need to pinpoint problem processes and their owners.
- Improve system performance by providing you with real-time information on which users and what processes are using the system.

Note: When monitoring process information, pay special attention to the following attribute values.

- Large values in the CPU Utilization column indicate a CPU-intensive process. These values result in a lower process priority. Small values indicate an I/O intensive process. These values result in a higher process priority.
- A value greater than 100 indicates a process is consuming a large amount of the CPU. If this value is high, check the Execution State column to see if the process is running, and the Time column to see how long the process has been running.

The CPU Utilization value is computer-dependent and also varies according to the version of the UNIX operating system running.

Virtual Size - Top Ten (bar chart)
The Virtual Size - Top Ten graph provides information about virtual memory process usage by the top ten processes.

The following workspaces are linked workspaces available for all the workspaces in the Process navigator item. They can be accessed by selecting the link indicator

next to a row of data in a view. Left-clicking a link indicator selects the default workspace associated with that link. Right-clicking a link indicator displays all linked workspaces that can be selected.

- Application for Process
- Child Processes
- Command for Process
- Process Resource
- Processes for Group Leader

RPC Performance workspace

The RPC Performance workspace is comprised of four views. This workspace contains the following views:

RPC Server Statistics (bar chart)
This chart graphically displays the following attributes:

- RPC Client Calls Rejected by Server
- RPC Server Times RPC Packet Unavailable
- RPC Server Packets Too Short
- RPC Server Packets with Malformed Header

RPC Client Statistics (bar chart)
This chart graphically displays the following attributes:

- RPC Client Calls Rejected by Server
- RPC Client Calls Retransmitted
- RPC Client Replies Not Matching Calls
- RPC Client Calls Timed Out

RPC Server Performance report (table view)
This table contains statistics attributes that refer to Remote Procedure Call (RPC) Server call rates and errors. The RPC Server attributes report on RPC calls made to the managed system.

RPC Client Performance report (table view)
This table contains statistics attributes that refer to Remote Procedure Call (RPC) Client call rates and errors. The RPC Client attributes report on calls made by the managed system to RPC Servers.

Solaris System CPU Workload workspace

The Solaris System CPU Workload workspace, accessed by right-clicking on the System Information navigator item, includes the Solaris System CPU Workload table chart, the Solaris System CPU Workload Statistics bar chart, and the Take Action view.

Solaris Zone Processes workspace

The Solaris Zone Processes workspace includes the Top CPU Time in zone, Top Virtual Size in zone, and Top Memory % used in zone views, accessed by right-clicking the Process navigator item.

Solaris Zones workspace

The Solaris Zones workspace, accessed by right-clicking on the System Information navigator item, includes the Solaris Zones and Take Action views.

System Details workspace

The System Details workspace, accessed by right-clicking on the System Information navigator item, includes System CPU, System Virtual Memory, Systems, System Load Average, System Workload, and System Cache and Buffers views.

System Information workspace

The System Information workspace is comprised of five views. This workspace contains the following views:

Virtual Memory Availability (pie chart)
This chart displays the following attributes:

- Virtual Memory Percent Used
- Virtual Memory Percent Available

Memory shortages can cause system performance problems. If system performance is poor, excessive page-outs and swapping activity can indicate memory problems. By viewing the monitored virtual memory data collected by IBM Tivoli Monitoring for UNIX on a remote system, you can perform the following tasks:

- Determine whether performance degradation is caused by a lack of virtual memory.
- View monitored data collected from remote systems either as a report or a chart.

Page Fault Statistics (bar chart)
This chart displays page fault attributes to provide an at-a-glance view of problems with virtual memory faults. The attributes displayed are:

- Page Faults
- Page Scan Rate
- Page Reclaims
- Pages Paged In
- Pages Paged Out
- Page Ins
- Page Outs

Total Real and Virtual Memory (bar chart)
The Total Real and Virtual Memory chart graphically displays vital information about system memory. The attributes displayed are:

- Total Real Memory
- Free Memory
- Total Virtual Memory

The chart gives you operating system and memory information on your UNIX systems. By viewing the monitored system data collected by IBM Tivoli Monitoring: UNIX OS Agent on remote systems, you can perform the following tasks:

- Improve system performance by helping you identify the configuration of your systems and check their current activity levels
- View monitored data collected from remote systems either as a report or as a chart.

CPU % (pie chart)
The CPU % chart helps you improve system CPU performance, and you

can use it to identify and monitor system CPU activity. The CPU % chart displays percentages of processor activity taking place on each monitored UNIX system. Use this report to perform the following tasks:

- Improve system CPU performance by helping you identify managed systems that consume large amounts of CPU time
- Increase system throughput by identifying user demands on CPUs, allowing you to allocate these demands among several CPUs on your system
- Identify managed systems with I/O bottlenecks caused by waits for CPU time
- Identify managed systems with CPU bottlenecks caused by high CPU utilization

Use the CPU % chart to check for problems, such as:

- Managed systems with high CPU utilization
- Imbalances between user and system CPU demands
- Long CPU waits caused by I/O bottlenecks

Note: There is a limit of up to 9 CPU per page. You can either modify the workspace or select a second page.

Load Averages (bar chart)

The Load Average chart gives you load average information on your UNIX systems. Load average refers to the average number of processes in the UNIX kernel run queue during an interval. By viewing the monitored system data collected by IBM Tivoli Monitoring: UNIX OS Agent on remote systems, you can improve system performance by helping you identify the configuration of your systems and check their current activity levels.

Top CPU-Memory %-VSize Details workspace

The Top CPU-Memory %-VSize Details workspace includes the Top CPU Time, Top Virtual Size, Top Memory % Used, and Take Action table views, accessed by right-clicking the Process navigator item.

UNIX workspace

The UNIX workspace, accessed by right-clicking the UNIX navigator tree item, is comprised of the following views:

Top CPU Time (bar chart)

Memory % - Top Ten (bar chart)

Disk Space Used % - Top Ten (bar chart)

Virtual Size - Top Ten (bar chart)

Virtual Memory Availability (bar chart)

UNIX Detail workspace

The UNIX Detail workspace, accessed by right-clicking the UNIX navigator tree item, is comprised of the following views:

Top CPU Time - Processes (table view)

Top Virtual Size (table view)

Top Memory % (table view)

Top Space Used % (table view)

System Virtual Memory (table view)

Users workspace

The Users workspace contains the following views:

Active Users (table view)
The Active Users attributes refer to user characteristics such as idle time, user name, location, and login time.

Defined Users (table view)
Use the Defined Users view to display information about the defined users, including logon user names, defined roles, whether the user accounts have been locked, expiration dates, and the number of incorrect logon attempts before the account is locked.

Chapter 4. Attributes reference

This chapter contains information about the following topics:

- Overview of attributes
- References for detailed information about attributes
- Descriptions of the attributes for each attribute group included in this monitoring agent
- Disk space requirements for historical data

About attributes

Attributes are the application properties that are measured and reported by the Monitoring Agent for UNIX OS, such as the amount of memory usage and the message ID.

Attributes are organized into groups according to their purpose. The attributes in a group can be used in the following two ways:

- Chart or table views

  Attributes are displayed in chart and table views. The chart and table views use queries to specify which attribute values to request from a monitoring agent. You use the Query editor to create a new query, modify an existing query, or apply filters and set styles to define the content and appearance of a view based on an existing query.
- Situations

  You use attributes to create situations that monitor the state of your operating system, database, or application. A situation describes a condition you want to test. When you start a situation, the Tivoli Enterprise Portal compares the values you have assigned to the situation attributes with the values collected by the Monitoring Agent for UNIX OS and registers an *event* if the condition is met. You are alerted to events by indicator icons that are displayed in the Navigator.

The Monitoring Agent for UNIX OS provides the following types of attributes:

single-instance
: Single-instance attributes are attributes that gather only one set of data. For example, the local time attributes are single-instance attributes because there is only one set of values for local time at any one time.

multiple-instance
: multiple-instance attributes are attributes that can gather multiple sets of data. For example, the Avg_Queue attribute is a multiple-instance attribute because it can return one set of data for each queue that exists on the system.

  You cannot use attributes from more than one multiple-instance attribute group in the same situation. Examples of multiple-instance attribute groups are Disk_Performance, System, and User.

Some of the attributes in this chapter are listed twice, with the second attribute having a "(Unicode)" designation after the attribute name. These Unicode attributes were created to provide access to globalized data.

More information about attributes

For more information about using attributes and attribute groups, see the *IBM Tivoli Monitoring User's Guide*.

For a list of the attributes groups, a list of the attributes in each attribute group, and descriptions of the attributes for this monitoring agent, refer to the Attribute groups and attributes section in this chapter.

Attribute groups and attributes

An attribute is a characteristic of a managed object (node). For example, Disk Name is an attribute for a disk, which is a managed object.

Use IBM Tivoli Monitoring: UNIX OS Agent attributes to build situations that monitor the performance of your UNIX network managed systems. When the values of the selected attributes in a situation exceed their threshold settings, the agent posts an alert to the Tivoli Enterprise Portal console notifying you of a problem.

Attributes groups contain System Name and Timestamp attributes. The following entries describe these attributes.

System Name
The managed system name. The form should be *hostname:agent_code*. Examples include spark:KUX or deux.raleigh.ibm.com:KUX. In workspace queries, this attribute should be set equal to the value $NODE$ in order to populate the workspace with data. This attribute is generally not included in situations, unless there is a need to customize the situation for a specific managed system.

Timestamp
The date and time the agent collects information as set on the monitored system. The timestamp for SCAN and STR functions is in the CYYMMDDHHMMSSmmm format (as in 1020315064501000 for 03/15/02 06:45:01) where:

C	= Century (0=20th)
YY	= Year
MM	= Month of the Year (01-12)
DD	= Day of the Month (01-31)
HH	= Hour, in 24-hour time (00-23)
MM	= Minute
SS	= Second
mmm	= Millisecond

Some of the attributes have the enumerations, Value Exceeds Maximum and Value Exceeds Minimum. The Tivoli Enterprise Monitoring Server allows only signed integers, so the maximum is 2147483647 and the minimum is -2147483648. If the agent has a value bigger or smaller than these, it is capped with these enumerations.

Workspaces and attributes

The IBM Tivoli Monitoring: UNIX OS Agent workspaces provide real-time information on many of the attributes. Each column in a table view corresponds to an attribute. And each data series in a chart corresponds to an attribute.

Groups of attributes

Each attribute belongs to an attribute group. The attribute group includes attributes that are related. Each attribute item stores data for a particular property of an attribute group.

The following are the attribute groups for IBM Tivoli Monitoring: UNIX OS Agent. The groups are collected in attribute tables that are designated in brackets [ ] after the group name.

- Agent Availability Management Status [KUXPASMGMT]
- Agent Active Runtime Status [KUXPASSTAT]
- AIX AMS [UNIXAMS]
- AIX Defined Users [UNIXDUSERS]
- AIX Devices [UNIXDEVIC]
- AIX LPAR [UNIXLPAR]
- AIX WPAR CPU [UNIXWPARCP]
- AIX WPAR File System [UNIXWPARFS]
- AIX WPAR Information [UNIXWPARIN]
- AIX WPAR Network [UNIXWPARNE]
- AIX WPAR Physical Memory [UNIXWPARPM]
- Alerts Table [KUXPASALRT]
- All Users Group [UNIXALLUSR]
- Configuration Information [KUXPASCAP]
- Disk Information [UNIXDISK]
- Disk Performance [UNIXDPERF]
- File Comparison Group [UNIXFILCMP]
- File Information [FILEINFO]
- File Pattern [UNIXFILPAT]
- Group [UNIXGROUP]
- IP Address [UNIXIPADDR]
- Machine Information [UNIXMACHINE]
- Network [UNIXNET]
- NFS/RPC Statistics [UNIXNFS]
- Ping Group [UNIXPING]
- Print Queue [UNXPRINTQ]
- Process [UNIXPS]
- SMP_CPU [UNIXCPU]
- Solaris Zones [UNIXSOLZON]
- System group [UNIXOS]
- TCP Statistics [UNIXTCP]
- UNIX Memory [UNIXMEM]
- User [UNIXUSER]

IBM Tivoli Monitoring provides other attribute groups that are available to all monitoring agents, for example Universal Time and Local Time. The attributes in these common attribute groups are documented in the Tivoli Enterprise Portal Help.

Platforms supported

Attributes display information on all of the platforms below unless noted otherwise in the attribute usage information.

- AIX
- HP-UX
- Solaris

Assignment of values

Not all UNIX systems display all UNIX attributes. For example, AIX systems do not display the CPU ID on which the process is running. If your system does not display a value for a certain attribute, you see **Not_Available** in fields relating to that attribute. A value of -1, or Not_Available, means that this information is not currently being collected for the UNIX platform on which your system is running. For example, HP-UX and Solaris systems do not display Busy Percent information.

A value of -2, or **Not_Collected**, is an indication that an error occurred.

A value of either 9,223,372,036,854,775,807, or **Value_Exceeds_Maximum**, and -9,223,372,036,854,775,808, or **Value_Exceeds_Minimum**, indicates that the value is too large and the Tivoli Enterprise Monitoring Server cannot handle the value without wrapping it. These values are then capped at either 9,223,372,036,854,775,807 or -9,223,372,036,854,775,808.

Cross referencing historical reports and attributes

Historical reports use a column header that identifies the attributes using an shorter character name. The historical column header is identified in capital letters surrounded by brackets [ ] under the attribute name. The historical data tables are identified in the same manner after the attribute group name.

Here is an example of an attribute:

Space_Used_Percent
Is the attribute name

[PCTSPCUSED]
Is the historical column header.

Here is an example of an attribute group:

File_Information Group [UNIXFILE]
Is the name of the attribute group and the name of the historical table.

Agent Availability Management Status attributes

Use Agent Availability Management Status attributes to view the current management status of an agent relative to Agent Management Services.

Agent Management Status The watched agent management status. The following values are valid: Unmanaged (0), Managed (1), Watchdog (2). A value of 'Managed' means that the agent is under the management of Agent Management Services. A

value of 'Unmanaged' means it is known, but that the agent is not under the management of Agent Management Services.

Agent Name The watched agent name.

Agent Type The watched agent type. The following values are valid: Unknown (0), ITM_Unix (1), Console (2), Windows_Service (3), Discover_ITM (4), Discover_Bin (5), Linux_Service (6), ITM_Windows (7).

Agent Version The version, release, and modification information for the agent.

Build Number The build number information for the agent.

Manager Type The enum defining the manager type. The following values are valid: Unknown (0), Not_Managed (1), Agent_Management Services (2), Watchdog (3), External (4). A value of 'Agent Management Services' means that Agent Management Services is responsible. A value of 'NotManaged' means that the agent is not under availability monitoring by any application. A value of 'Externally' means that some other application besides Agent Management Services is responsible for availability monitoring of the agent, for example Tivoli System Automation or Windows service control manager.

Operating System The operating system identification. The following values are valid: Unknown (0), Windows (1), Linux (2), UNIX (3).

Server Name The origin node of the collecting agent.

Service Name The Windows or Linux service name of the agent.

Timestamp The date and time the Tivoli Enterprise Monitoring Server samples the data.

Agent Active Runtime Status attributes

Use the Agent Active Runtime Status attributes to view the current availability status of an agent: Running, Not present, Unknown, Stopped, Manually Stopped. You can view the frequency at which the agent's availability and runtime properties are queried and also the agent's Daily Restart Count.

Agent Availability Status The watched agent availability status. The following values are valid: Unknown (0), Not_found (1), Stopped (2), Start_Pending (3), Running (4), Manually_Stopped (5), Stop_Pending (6), Not_configured (7). For agents that have an Availability Status of 'Running', use the attribute group to see runtime properties of the agent such as its Process ID and Thread Count.

Agent Host Name The host name of the agent.

Agent Name The watched agent name.

Agent Type The watched agent type. The following values are valid: Unknown (0), ITM_Unix (1), Console (2), Win_Service (3), Discover_ITM (4), Discover_Bin (5), Linux_Service (6), ITM_Windows (7).

Check Frequency The frequency to check status in seconds.

Command Line The command line.

Daily Restart Count The restarts within a period of a day.

Instance Name The instance name of the running managed IBM Tivoli Monitoring agent.

IP Address The IP address of the agent.

Last Health Check The last health check timestamp.

Number of Threads The thread count.

Operating System The operating system identification. The following values are valid: Unknown (0), Windows (1), Linux (2), UNIX (3).

Page Faults Per Second The page faults per second.

Parent Process ID The parent process identification.

Process ID The process ID.

Process Name The process name.

Process System CPU (Percent) The system CPU.

Process User CPU (Percent) The user CPU time.

Resident Size The size of the resident process.

Server Name The origin node of the collecting agent.

Timestamp The date and time the Tivoli Enterprise Monitoring Server samples the data.

Total Size (Pages) The total memory size in pages.

User Name The user name of the running managed agent.

AIX AMS attributes

Use the AIX AMS attributes to view information about the Active Memory Sharing (AMS) pool.

AMS Mem Loaned AMS logical memory loaned to the hypervisor.

AMS Memory Ent InUse AMS memory entitlement of the partition in use (MB).

AMS Memory Entitlement AMS memory entitlement of the partition (MB).

AMS Mode Indicates whether the LPAR is in AMS shared or dedicated mode.

AMS Physical Mem Physical memory supporting AMS logical memory for the partition.

AMS Pool ID The pool ID associated with the LPAR. All LPARs in AMS mode will have a pool ID of 0 until multiple pools are supported.

AMS Pool Size AMS Memory pool size in GB.

Hypervisor Page Ins Number of hypervisor page-ins.

Hypervisor Page Ins Time Time spent waiting for hypervisor page-ins in nanoseconds.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

AIX Defined Users attributes

Use the AIX Defined Users attributes to view information about the defined users, including logon user names, defined roles, whether the user accounts have been locked, expiration dates, and the number of incorrect logon attempts. Note that the attribute values include Not Available and Not Collected, when applicable. Note that the data collection for this attribute group is by default disabled for performance reasons. You must set KUX_DEFINED_USERS=True in the `ux.ini` file to enable it.

Account Locked The indicator of whether or not the user account has been locked.

Expires The expiration date of this user ID.

Login Retries The number of incorrect logon attempts before the user ID is locked.

Roles The roles defined for this user ID.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

User Name The logon user name.

AIX Devices attributes

Use the AIX Devices attributes to view information about network and storage device status.

Class The class of the device.

Name The name of the device.

Parent The parent device name.

System Name The managed system name.

State The device status.

Timestamp The date and time the agent collects information as set on the monitored system.

Type The device type.

AIX LPAR attributes

Use the AIX LPAR attributes to view information about the logical partition (LPAR).

Available CPU Units in Pool The number of physical CPU units that are available for allocation from the shared pool.

Available CPUs in Pool The number of CPUs that are available for allocation.

Busy Pct The logical busy time percentage.

Capacity Weight The relative weight between 0 and 255 that is used to determine how much extra CPU capacity this LPAR is to receive.

Capped Mode The capped Logical Partition mode (uncapped or capped).

CPU Entitlement The entitled processor capacity for the partition.

CPU Pool ID The ID of the Shared Processor Pool.

Donated Busy Cycles Pct The percentage of physical processor that is used by donating busy cycles, for dedicated partitions only. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

Donated Idle Cycles Pct The percentage of physical processor that is used by explicitly donated idle cycles, for dedicated partitions only. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

Donation Enablement The status of the willingness of this LPAR to allow unused CPU cycles to be used by other LPARs. Note: the value -1 indicates Not Collected, 0 indicates disabled, 1 indicates capable, and 2 indicates enabled.

Entitlement The number of entitlement units assigned to this LPAR.

Entitlement Pct The entitlement as a percentage.

Entitlement Used Pct The percentage of the given CPU Entitlement being used by this LPAR.

Hostname The host name of the LPAR.

Hypervisor Calls The number of hypervisor calls made during the monitoring period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Last Machine ID The previous frame hardware ID of this LPAR before it was migrated to the current frame. This attribute is cached for 5 minutes after the LPAR migration is detected to allow a situation to discover and report that the LPAR had been migrated.

LPAR Name The name of the logical partition.

LPAR Number The LPAR identification number assigned to this LPAR.

Machine ID The frame hardware ID to which this LPAR belongs.

Max CPU Cap Used Pct The percentage of maximum physical CPU available to this LPAR that was actually used. For capped LPARs, this value is the same as CPU Phys Ent Pct.

Maximum Pool Capacity The maximum pool capacity.

Number of Logical CPUs The number of current online logical CPUs.

Number of Physical CPUs The number of active licensed physical CPUs.

Number of Physical CPUs in Shared Pool The number of physical CPUs in the shared pool.

Number of Virtual CPUs The number of current online virtual CPUs.

Phantom Interrupts The number of phantom interrupts.

Phys Busy Pct The physical busy time of a full processor percentage.

Physical CPU Size of Shared Pool The number of physical CPU units in the shared pool.

Physical CPU Units Used The number of physical CPU units consumed by this LPAR.

Pool Entitlement The entitled capacity of the pool.

Shared Mode The shared Logical Partition mode (dedicated or shared).

SMT Mode The simultaneous multi-threading mode (off or on).

SMT Threads The number of threads per CPU.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Total Used Pct The percentage of the Total System CPU being used by this LPAR.

Unallocated CPU In Pool The unallocated capacity available in the shared pool.

Uptime The period of time this LPAR has been operational.

Virt Context CPU Switches per Sec The virtual CPU context switches per second.

AIX WPAR CPU attributes

Use the AIX WPAR CPU attributes to view CPU usage information for the workload partitions (WPARs).

CPU Consumption Limit The maximum number of physical processors that a WPAR is allowed to consume.

LPAR CPU Consumed Pct The percentage of CPU entitlement of the LPAR consumed by the WPAR.

LPAR Entitlement The number of entitlement units assigned to this LPAR.

Num CPUs Consumed The number of physical processors consumed by the WPAR when the LPAR is in shared processor mode.

RC CPU Limits Hard Max The maximum percentage CPU that a WPAR can have even if there is no contention for CPU.

System CPU Pct The time this WPAR spent running in CPU kernel mode percentage.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

User CPU Pct The time this WPAR spent running in CPU user mode percentage.

WPAR CPU Consumed Pct The percentage of its CPU share consumed by the WPAR.

WPAR Name The name of the workload partition (WPAR).

AIX WPAR File System attributes

Use the AIX WPAR File System attributes to view file system information of the workload partitions (WPARs).

Device Name The name of the mounted file system.

Mount Options The specified mount options.

Mount Point The file system mount point.

Node Name The name of the remote node.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

VFS Type The virtual file system type.

WPAR Name The name of the workload partition (WPAR).

AIX WPAR Information attributes

Use the AIX WPAR Information attributes to view the general configuration parameters of the workload partitions (WPARs).

Admin Operation The administrative operation being performed.

Admin Process ID The PID of the administrative operation being performed.

Admin Start Time The time when the administrative operation started.

Autostart Indicates whether the WPAR starts on a restart.

Checkpointable Indicates whether the WPAR can be checkpointed or not.

Home The home directory for the WPAR.

Hostname The host name for the WPAR.

IP Address The IP or Network address of the network interface.

Owner The user ID that owns the WPAR.

RC CPU Limits Hard Max The maximum percentage of CPU that a WPAR can have even if there is no contention for CPU.

RC CPU Limits Min The minimum percentage of CPU that is guaranteed to the WPAR.

RC CPU Limits Soft Max The maximum percentage of CPU that a WPAR can have when there is a contention for CPU.

RC CPU Shares The number of CPU shares for this WPAR.

RC Is Active Indicates whether the resource controls are active or not.

RC Max Processes The total number of processes allowed in the WPAR.

RC Max Threads The total number of threads allowed in the WPAR.

RC Memory Limits Hard Max The maximum percentage of memory that a WPAR can have even if there is no contention for memory.

RC Memory Limits Min The minimum percentage of memory that is guaranteed to the WPAR.

RC Memory Limits Soft Max The maximum percentage of memory that a WPAR can have when there is a contention for memory.

RC Memory Shares The number of memory shares for this WPAR.

RC per Process VM Limit The maximum amount of virtual memory that a process in the WPAR can consume.

RC RSet The name of the configured Resource Set.

Shares usr Dir Indicates whether the WPAR shares its /usr file system with the LPAR.

State The current state of the WPAR, including Defined, Loaded, Active, Frozen, Paused, Transitional, and Broken.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Type The type of the WPAR, including application and system.

WPAR Application Path The full path of the executable file to run inside the Application WPAR.

WPAR Name The name of the workload partition (WPAR).

AIX WPAR Network attributes

Use the AIX WPAR Network attributes to view network-related information of the workload partitions (WPARs).

Broadcast IP The broadcast IP address.

Interface Name The name of the network interface.

IP Address The IP or network address of the network interface.

Network Mask The internet network mask.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

WPAR Name The name of the workload partition (WPAR).

AIX WPAR Physical Memory attributes

Use the AIX WPAR Physical Memory attributes to view general configuration parameters of the workload partitions (WPARs).

Free Memory MB The amount of free (unallocated) memory in the WPAR in MB.

Free Memory Pct The percentage of free memory in the WPAR.

LPAR Memory Size MB The total amount of physical memory available to the LPAR in MB.

LPAR Memory Used Pct The percentage of LPAR memory used by the WPAR.

Memory Size MB The total amount of physical memory available to the WPAR in MB.

RC Memory Limits Hard Max The maximum percentage of memory that a WPAR can have even if there is no contention for memory.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Used Memory MB The amount of used (allocated) memory in the WPAR in MB.

Used Memory Pct The percentage of used memory in the WPAR.

WPAR Name The name of the workload partition (WPAR).

Alerts Table attributes

Use the Alerts Table attributes to view exceptional Critical, Warning, or Informational alerts sent by Agent Management Services. You can view these for 24 hours or until the OS agent is recycled. These events have to do with the operation of Agent Management Services or conditions affecting its ability to manage agents. The following alerts are included:

- Availability policy removed
- Agent abnormally stopped
- Agent restart failed
- Agent exceeded restart count
- Agent not found.
- Agent overutilizing memory
- Agent overutilizing CPU
- Managed agent removed from system
- Unmanaged agent removed from system
- Agent start failed
- Agent status check script failed

Agent Name The watched agent name.

Agent Status The agent status. The following values are valid: Unknown (0), Not_found (1), Stopped (2), Start_Pending (3), Running (4), Manually_Stopped (5), Stop_Pending (6), Not_configured (7).

Agent Type The watched agent type. The following values are valid: Unknown (0), ITM_Unix (1), Console (2), Windows_Service (3), Discover_ITM (4), Discover_Bin (5), Linux_Service (6), ITM_Windows (7).

Alert Details The alert message details.

Alert Message The alert message. The following values are valid: Availability_policy_removed (1), Managed_agent_removed_from_system (2), Unmanaged_agent_removed_from_system (3), Agent_abnormally_stopped (4), Agent_exceeded_restart_count (5), Agent_restart_failed (6), Agent_overutilizing_memory (7), Agent_overutilizing_CPU (8), Agent_manual_stop_failed (9), Agent_Management_Services_watchdog_not_reliable (11).

Operating System The operating system identification. The following values are valid: Unknown (0), Windows (1), Linux (2), UNIX (3).

Process ID The process ID.

Process Name The process name.

Server Name The origin node of the collecting agent.

Timestamp The date and time the Tivoli Enterprise Monitoring Server samples the data. This information is displayed in the standard 16-character date/time format (CYYMMDDHHMMSSmmm), where:

C	Century (0 for 20th, 1 for 21st)
YY	Year
MM	Month
DD	Day
HH	Hour
MM	Minute
SS	Second
mmm	Millisecond

Use simple text strings as described above. For example, 1101009130500000 expresses October 9, 2010, 1:05:00 pm.

All Users attributes

The All Users attributes refer to user characteristics such as user ID and user sessions.

Duplicate User Name True if the user name is listed more than once in /etc/passwd. True if no password is assigned to the user. The following values are valid: Not Collected, Not Available, False and True.

Name The full name of a user.

No Password True if no password is assigned to the user. The following values are valid: Unknown, False, Not Collected, and True.

System Name The managed system name. The form should be *hostname:agent_code*.

Examples include `spark:KUX` or `deux.raleigh.ibm.com:KUX`.

In workspace queries, this attribute should be set equal to the value $NODE$ in order to populate the workspace with data. This attribute is generally not included in situations, unless there is a need to customize the situation for a specific managed system.

Timestamp The date and time the agent collects information as set on the monitored system.

User ID The numeric ID the system assigned to a user. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

User Sessions The number of login sessions this user currently has established. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Configuration Information attributes

Use Configuration Information attributes to monitor agent configuration such as Memory Threshold and Operating System.

Agent Name The sub agent name.

Agent Path The fully qualified path to the agent.

Agent Type The watched agent type. The following values are valid: Unknown (0), ITM_Unix (1), Console (2), Windows_Service (3), Discover_ITM (4), Discover_Bin (5), Linux_Service (6), ITM_Windows (7).

Check Frequency The frequency to check status in seconds.

Configuration Script The agent configuration script.

% CPU Threshold The amount of CPU the agent process is allowed to consume before Agent Management Services restarts it.

Dependencies Any agents, monitored by Agent Management Services, that must be started before the target agent can be started.

Kernel Release The version of Linux kernel that a particular CAP file stanza applies to, if that level of granularity is needed.

Manager Type The enum defining the manager type. The following values are valid: Unknown (0), Not_Managed (1), Agent_Management Services (2), Watchdog (3), External (4).

Maximum Daily Restarts The maximum number of restarts allowed. The clock begins at midnight.

Memory Threshold The amount of working set memory the agent process is allowed to consume before Agent Management Services restarts it.

Memory Unit The maximum memory allowed units. The following values are valid: Bytes (0), KB (1), MB (2), GB (3).

Operating System The operating system that a particular CAP file stanza applies to. The following values are valid: Unknown (0), Windows (1), Linux (2), and UNIX (3).

Operating System Name The operating system name that a particular CAP file stanza applies to, if that level of granularity is required.

Operating System Version The operating system version that a particular CAP file stanza applies to, if that level of granularity is required.

PAS_ID The internal ID assigned to an agent by Agent Management Services.

Policy File Timestamp The date and time of CAP file.

Process Name The process name of the managed agent.

Server Name The origin node of the collecting agent.

Service Name The Windows or Linux service name of an agent.

Startup Script The agent startup script.

Status Script The agent status script.

Stop Script The agent stop script.

Timestamp The date and time the Tivoli Enterprise Monitoring Server samples the data. This information is displayed in the standard 16-character date/time format (CYYMMDDHHMMSSmmm), where:

C	Century (0 for 20th, 1 for 21st)
YY	Year
MM	Month
DD	Day
HH	Hour
MM	Minute
SS	Second
mmm	Millisecond

Use simple text strings as described above. For example, 1101009130500000 expresses October 9, 2010, 1:05:00 pm.

Disk attributes

The Disk attributes refer to disk characteristics such as inode size, inodes used, mount point, and space available. Only mounted file systems are monitored by this agent. Entries for the Disk Information table (UNIXDISK) reports file systems of the following types:

- AIX systems
  - GPFS - General Parallel File System
  - JFS - Journaled File System
  - JFS2 - Enhanced Journaled File System
  - NFS - Network File System
  - Veritas VxFS
- Solaris systems
  - ufs - UNIX File System
  - tmpfs - Temp File System
  - vxfs - Veritas File System
  - bfs -Boot File System
  - lofs - loopback File System
  - zfs - Zettabyte File System
  - NFS - Network File System
- HPUX systems

– All File Systems

For UNIX OS Agents running on AIX machines, the keywork KBB_SHOW_CUSTOMFS enables monitoring of custom filesystems that are not available in the list. Typically, the values range between 8 - 31 (defined as part of /usr/include/sys/vmount.h AIX file). To monitor MMFS (Multimedia File System) on AIX, add KBB_SHOW_CUSTOMFS=8 to the ux.ini file.

File System Status The availability status of the remote file system (NFS). Note: the value 2 indicates Up, 1 indicates Down, 0 indicates Not Available, and -2 indicates Not Collected.

File System Type The filesystem type name, for example jfs and gpfs. Values that include up to eight letters or numbers are valid. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Any query, applied to an agent prior to V6.2.3, containing this attribute displays a blank in the File System Type column.

Inodes Available Percent The percentage of inode space currently not in use. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Inodes Free The number of inodes currently available on your file system. Use this attribute to avoid a pending crisis. Corrective action might include freeing up unneeded space or deleting temporary files. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Inodes Free (Superseded) The number of inodes currently available on your file system. Use this attribute to avoid a pending crisis. Corrective action might include freeing up unneeded space or deleting temporary files. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Example: Disk inode space is usually a high number, so the percentage of inode space to disk space must also be a high number. If the value for Inodes Free is less than 100, this is a critical condition. Notify your system administrator immediately. If the value for Inodes Percent is above 30%, delay taking corrective action.

Inodes Used The number of inodes currently allocated to files on the file system. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Inodes Used (Superseded) The number of inodes currently allocated to files on the file system. This value equals the Inode Size value minus the Inodes Free value. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Inodes Used Percent The percentage of inode space currently allocated to files. Valid entries are whole numbers up to 100, such as 85 for 85%. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Mount Point The path name of the directory to which a file system is mounted. This is the virtual name for the directory. Valid entries are up to 32 letters or numbers representing a directory path.

Mount Point (Unicode) The path name of the directory to which a file system is mounted. This is the virtual name for the directory.

Name The name of the physical disk partition where the file system is mounted. This is the physical location of the disk. Valid entries are up to 32 letters or numbers.

Name (Unicode) The name of the physical disk partition where the file system is mounted. This is the physical location of the disk.

Size (KBytes) The total size of a file system, expressed in kilobytes with 64 bit precision. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Size (KBytes) (Superseded) The total size of a file system, expressed in kilobytes. For example, 1000000 represents one gigabyte. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Size (GigaBytes) The total size of a file system, expressed in gigabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Size (GB) (Superseded) The total size of a file system, expressed in gigabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Size (MegaBytes) The total size of a file system, expressed in megabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Size (MB) (Superseded) The total size of a file system, expressed in megabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Space Available (KBytes) The amount of unused space currently available to non-superusers on a file system, expressed in kilobytes. For example, 40000 represents 40 megabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Space Available (KB) (Superseded) The amount of unused space currently available to non-superusers on a file system, expressed in kilobytes. For example, 40000 represents 40 megabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

This disk space does not include any space which is reserved for supersuser. A low value in this column, relative to the disk size, alerts you to critical disk space conditions.

If this value is low for one or more file systems, relative to the disk size, you might need to evaluate reconfiguring the file system to distribute the files more evenly across disks.

Space Available (GigaBytes) The amount of disk space currently available to non-superusers on a file system, expressed in gigabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9,223,372,036,854,775,807 indicates Value_Exceeds_Maximum, and -9,223,372,036,854,775,808 indicates Value_Exceeds_Minimum.

Space Available (GigaBytes) (Superseded) The amount of disk space currently available to non-superusers on a file system, expressed in gigabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Space Available (MegaBytes) The amount of disk space currently available to non-superusers on a file system, expressed in megabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Space Available (MegaBytes) (Superseded) The amount of disk space currently available to non-superusers on a file system, expressed in megabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Space Available Percent The percentage of space available. Valid entries are whole numbers up to 100, such as 10 for 10%. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Space Used (KBytes) The amount of disk space currently in use on a file system, expressed in kilobytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Space Used (KB) (Superseded) The amount of disk space currently in use on a file system, expressed in kilobytes. Valid entries For example, 5000 represents five megabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Space Used (GigaBytes) The amount of disk space currently in use on a file system, expressed in gigabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Space Used (GigaBytes) (Superseded) The amount of disk space currently in use on a file system, expressed in gigabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Space Used (MegaBytes) The amount of disk space currently in use on a file system, expressed in megabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Space Used (MegaBytes) (Superseded) The amount of disk space currently in use on a file system, expressed in megabytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Space Used Percent The space currently used on the file system, expressed as a percentage of the sum of used and available space. The Space Used Percent reflects the percentage of disk space which is available to non-superusers. A high value in this column alerts you to critical disk space conditions. Valid entries are whole numbers up to 100, such as 80 for 80%. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Total Inodes The number of inodes allocated on a file system. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Total Inodes (Superseded) The number of inodes allocated on a file system. For example, a value of 163817 indicates that the number of inodes allocated is 163,817. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Use this attribute when a file system needs additional or fewer inodes assigned to it. Viewing the current number of inodes assigned helps you determine the number of inodes you need to add or subtract to optimize performance in your system.

Volume Group Name (AIX) The name of the volume group. Valid entries are up to 96 letters or numbers.

Disk Performance attributes

The Disk Performance attributes refer to disk operations such as data transfer rates, average waiting times, and percentage busy.

% Disk Read Time The percentage of elapsed time that the disk drive was busy servicing read requests over the previous 30-second interval. Note: the value -1 indicates Not Available and -2 indicates Not Collected. HP-UX systems and AIX systems do not measure this attribute and a value of -1 will be reported.

Avg Disk Bytes Xfer The number of bytes per disk transfer averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Queue Average number of disk requests outstanding during the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Service Queue Size (AIX) Average service queue size in an AIX environment. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Service Time Average amount of disk time used in milliseconds over the sampling period. For example, in terms of a bank teller queue, it is the time spent at the teller's window. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Wait Average time waiting for disk access expressed in milliseconds. For example, in terms of a bank teller queue, it is the time from when you first join the queue until you advance to the teller window to be serviced. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Busy Percent The percentage of time a disk has been busy transferring data. Valid entries are whole numbers up to 100, such as 35 for 35%. The Busy Percent value lets you check whether a process is I/O bound. Values greater than 30% usually indicate excessive paging out to disk, or that a process is I/O bound. If the Busy Percent value is high (greater than 30%) and CPU utilization is also high (greater than 80%), your system is probably overloaded and experiencing degradation of performance. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Disk Name The name of the physical disk which might be partitioned. Valid entries are up to 32 letters or numbers.

Disk Name (Unicode) The name of the physical disk which might be partitioned.

Disk Read Bytes per Sec The number of bytes per second transferred from the disk during read operations over the previous 30-second interval. HP-UX systems do not measure this attribute and a value of -1 will be reported. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Disk Reads per Sec The number of read operations on the disk per second averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. HP-UX systems and AIX systems do not measure this attribute and a value of -1 will be reported.

Disk Write Bytes per Sec The number of bytes per second transferred to the disk during write operations over the previous 30-second interval. HP-UX systems do not measure this attribute and a value of -1 will be reported. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Disk Writes per Sec The number of write operations on the disk per second averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. HP-UX systems and AIX systems do not measure this attribute and a value of -1 will be reported.

Parent (AIX) The parent device name. Valid entries are up to 96 letters or numbers. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Percent Disk Write Time The percentage of elapsed time that the disk drive was busy servicing write requests over the previous 30-second interval. Note: the value -1 indicates Not Available and -2 indicates Not Collected. HP-UX systems and AIX systems do not measure this attribute and a value of -1 will be reported.

Service Queue Full per Sec (AIX) The number of times, per second, that the service queue becomes full (the disk does not accept any more service requests). Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Transfer Rate The number of data transfers per second during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Transferred Bytes (KBytes) The total number of kilobytes that have been transferred during the recording interval. The Transferred Bytes count is one indicator of how fast your disk is moving data. It does not account for variables, such as disk format and efficiency of space usage, that also affect the speed of data transfer. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Transfers KB per Sec (AIX) The amount of data transferred (read or written) to the drive in KBs per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Type (AIX) The type of device. Valid entries are up to 96 letters or numbers.

File Comparison Group attributes

The File Comparison Group attributes refer to file characteristics such as file compare result and file name.

File Compare Option Options that affect how the comparison is performed. The following values are valid: Plain, Ignore Whitespace, Ignore Case, Ignore Case Whitespace, Binary Not Available, and Not Collected.

File Compare Result The result of the file comparison between File_Name_1 and File_Name_2. The following values are valid: Same, Different, Not Available, and Not Collected.

File Name 1 Fully-qualified file name of one of the files to be compared.

File Name 2 Fully-qualified file name of the other of the files to be compared.

System Name The managed system name. The form should be *hostname:agent_code*.

Examples include spark:KUX or deux.raleigh.ibm.com:KUX.

In workspace queries, this attribute should be set equal to the value $NODE$ in order to populate the workspace with data. This attribute is generally not included in situations, unless there is a need to customize the situation for a specific managed system.

Timestamp The date and time the agent collects information as set on the monitored system.

File Information attributes

The File Information attributes refer to file and directory characteristics such as name, size, owner, access rights, and links. File Information is a multiple-instance attribute group.

Note: This attribute group is not historically collected.

Access This attribute defines a four-digit octal number representing the access rights for a file. You specify access rights using a four-digit number representing the permissions associated with a file. Each digit is the decimal equivalent of a binary three-bit string. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Valid entries are Numeric strings in the range 0000 to 7777, from left to right, the digits have the following meaning:

1st Digit	Determines whether, on execution, the file takes on the ID of the user or group that owns the file. This permission assignment applies to users that neither own the file they are trying to run, nor belong to the group that owns the file.
2nd Digit	Determines the access permissions of the user that owns the file.
3rd Digit	Determines the access permissions of the group that owns the file.
4th Digit	Determines the access permissions for other users.
From right to left, the bits for the first digit have the following meanings:	
1st bit	The meaning if this bit depends on the type of UNIX operating system you are monitoring.
2nd bit	If the value of this bit is 1, the system runs the file with the group ID of the group that owns the file. If the value of this bit is 0, the system runs the file with the group ID of the user that ran the file. If the file is a directory and this bit is 1, all files created in that directory inherit the group ID of that directory.

Checksum The value of the checksum.

Checksum Algorithm Only used in situations in conjunction with the Checksum. Note: the -1 value indicates Not_Available, and the -2 value indicates Not_Collected.

File The name of a file or directory. If the file is a symbolic link, the link name is shown in the Link Name attribute. Valid entries: simple text string, alphanumeric with a maximum length 256 characters.

File Content Changed A numeric indicator that the content of a file has changed. It is equivalent to noting a change in checksum between two samples. The following values are valid: No, Yes, and Not Available.

File Size (Bytes) The size, in bytes, of a file. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

File Size (Bytes) (Superseded) The size, in bytes, of a file. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

File Size (MB) The size, in megabytes, of a file. Since the attribute does not have decimal precision, previously any file with a size less than "1" would be displayed as "0". As of IBM Tivoli Monitoring v6.1 Fix Pack 5, the size of the file is rounded up to the nearest whole number. For example, 1.2 would display as 2, and 0.1 would display as 1. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

File Size (MB) (Superseded) The size, in megabytes, of a file. Since the attribute does not have decimal precision, previously any file with a size less than "1" would be displayed as "0". As of IBM Tivoli Monitoring v6.1 Fix Pack 5, the size of the file is rounded up to the nearest whole number. For example, 1.2 would display as 2, and 0.1 would display as 1. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

File (Unicode) The name of a file or directory. If the file is a symbolic link, the link name is shown in the Link Name attribute.

Group The name of the logical group to which a file owner belongs. Valid entries: simple text string, alphanumeric with a maximum length 16 characters

Group (Unicode) The name of the logical group to which a file owner belongs.

Last Accessed Time The date and time of the last file access. Valid entries are in the CYYMMDDHHMMSSmmm format, where:

C	= Century (0=20th)
YY	= Year
MM	= Month of the Year (01-12)
DD	= Day of the Month (01-31)
HH	= Hour, in 24-hour time (00-23)
MM	= Minute
SS	= Second
mmm	= Millisecond

Last Attr Chg Time The time when the file attributes were last modified.

Last Changed Time The date and time of the last change to a file. Valid entries are in the format CYYMMDDHHMMSSmmm, where:

C	= Century (0=20th)
YY	= Year
MM	= Month of the Year (01-12)
DD	= Day of the Month (01-31)
HH	= Hour, in 24-hour time (00-23)
MM	= Minute
SS	= Second
mmm	= Millisecond

Links The number of links to a file. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Link Name The name of the file for which this file is a symbolic link. If this field is blank, the file is not a link. Note that V120 agents do not place data in this field. Valid entries are simple text string, alphanumeric with a maximum length 256 characters.

Link Name (Unicode) The name of the file for which this file is a symbolic link. If this field is blank, the file is not a link. Note that V120 agents do not place data in this field.

Mode The string representation of the file mode.

Owner The name of the file owner. Valid entries are simple text string, alphanumeric with a maximum length 16 characters.

Owner (Unicode) The name of the file owner. Valid entries are simple text string, with a maximum length 768 bytes.

Path The full path containing a particular file or directory. Valid entries are a text string representing the full path of the file, alphanumeric with a maximum length 256 characters. SCAN is not supported for this attribute.

Path (Unicode) The full path containing a particular file or directory. SCAN is not supported for this attribute.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Type The type of file. The following values are valid:

Dir	= Directory
DirLink	= Directory Link
File	= File
FileLink	= File Link
Sock	= Socket
Link	= Link

Spec	= Special
Unknown	= Unknown

The maximum length is 12 characters.

File Pattern attributes

The File Pattern attributes refer to file and match characteristics such as match count and match pattern.

File Name Fully qualified file name which will be searched for lines matching a pattern.

Match Count The number of matches for the specified pattern in the specified file. Note: the -1 value indicates Not_Available, and the -2 value indicates Not_Collected.

Match Option Options that affect how the search is performed. The following values are valid: Normal, Ignore Case, Inverse Search, Match Whole Words, Not Available, and Not Collected.

Match Pattern The grep regular expression used to search for matching lines in File_Name.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Group attributes

The Group attributes refer to group characteristics.

Duplicated Group Name True if the group name is listed more than once in /etc/group. The following values are valid: False and True.

Group ID The ID of this group. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, the value 2147483647 indicates Value_Exceeds_Maximum, and the value Value_Exceeds_Minimum=-2147483648.

Group Name The name of the group.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

IP Address attributes

The IP Address attributes refer to IP address characteristics.

DNS Name The Domain Name Server (DNS) entry associated with the IP network address. Note: No_DNS_Entry is a valid value.

IP Address An IP address associated with the network interface.

IP Version An indicator as to whether the IP address is version 4 or version 6. The following values are valid: Not Available (-1), Not Collected (-2), IPv4 (4) and IPv6 (6).

Network Interface Name The name of the network interface.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Machine Information attributes

The Machine Information attribute group contains various items required by other Tivoli products. They include system hardware information.

Hardware Manufacturer The manufacturer of the hardware on which the agent is running.

Hardware Model The specific hardware model underlying the monitored operating system.

Machine Serial Number The serial number of the computer.

Name The host name for the computer.

Number of Physical Processors The number of physical processors on the computer. This number excludes secondary processor contexts, but might include virtual processors in some virtual environments. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processor Megahertz The detected processor speed for the targeted operating system. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System Board UUID The Universally Unique Identifier burned in to the system board. The following value is valid: Unknown (UNKNOWN).

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Virtual Machine Identifier The serial number or name of the operating system instance. It usually corresponds to a name identifying a virtual machine, but it can also identify an nPar or vPar on HP-UX, or the default OS.

Network attributes

The Network attributes refer to network characteristics such as received count, sent count, network interface name, and interface status.

Avg Input Packet Rate 1 Minute The average number of packets received on all network interfaces over last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Input Packet Rate 5 Minutes The average number of packets received on all network interfaces over last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Input Packet Rate 15 Minutes The average number of packets received on all network interfaces over last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Input Packet Rate 60 Minutes The average number of packets received on all network interfaces over last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Input Errors Rate 1 Minute The average number of packets with errors received on all network interfaces over last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Input Errors Rate 5 Minutes The average number of packets with errors received on all network interfaces over last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Input Errors Rate 15 Minutes The average number of packets with errors received on all network interfaces over last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Input Errors Rate 60 Minutes The average number of packets with errors received on all network interfaces over last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Output Packets Rate 1 Minute The average number of packets transmitted on all network interfaces over last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2

indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Output Packets Rate 5 Minutes The average number of packets transmitted on all network interfaces over last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Output Packets Rate 15 Minutes The average number of packets transmitted on all network interfaces over last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Output Packets Rate 60 Minutes The average number of packets transmitted on all network interfaces over last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Output Errors Rate 1 Minute The average number of packets transmission errors on all network interfaces over last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Output Errors Rate 5 Minutes The average number of packets transmission errors on all network interfaces over last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Output Errors Rate 15 Minutes The average number of packets transmission errors on all network interfaces over last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Output Errors Rate 60 Minutes The average number of packets transmission errors on all network interfaces over last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Bandwidth Utilization Percent (AIX) The percentage of physical network adapter bandwidth utilized. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Collisions The number of times during the monitoring interval that a packet transmitted by the network interface collided with another packet. This occurs when another interface on the same local network transmits a packet at nearly the same time. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine if a network interface has

an unacceptable number of packet collisions. Packet collisions cause the interface to retransmit the packet. With this increased traffic, the likelihood of future collisions increases. This can result in a steady increase of network traffic to critical levels.

Collisions Rate Avg 1 Minute The average number of collisions on all network interfaces over last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Collisions Rate Avg 5 Minutes The average number of collisions on all network interfaces over last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Collisions Rate Avg 15 Minutes The average number of collisions on all network interfaces over last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Collisions Rate Avg 60 Minutes The average number of collisions on all network interfaces over last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Hostname IPv4 DNS Name The Domain Name Server (DNS) entry associated with the first IP address of the network interface. The following values are valid: Sum_of_all_interface_metrics and Not_Resolvable.

Input Errors The number of packets with errors received by the interface during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Input Packet Errors Percent Portion of packets received over the previous 30-second interval with errors. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Interface DNS Name The Domain Name Server (DNS) entry associated with the first IP address of the network interface. The following values are valid: `Sum of all interface metrics` and `Not_Resolvable`. Valid entries are alphanumeric strings with a maximum 32 characters.

Example:

www.company.com indicates that the DNS resolves the name www.company.com to mean the IP address for the interface.

Interface IP Address The Internet Protocol (IP) address of the network interface. A gateway system has more than one interface, each with a separate address. Valid entries: Internet protocol addresses in the form a.b.c.d. where a, b, c, and d are integers in the range 0 to 255.

Example 197.128.55.55 indicates the network interface uses the IP address 197.128.55.55. Multiple addresses assigned to a single network interface are not all displayed from Tivoli Enterprise Portal.

Interface Status This attribute indicates if a network interface is currently available. Valid entries for each Network interface:

UP	Indicates the interface is in service
DOWN	Indicates the interface is not in service
Up Not Running	Indicates the interface is in service but not running

These values are case-sensitive.

Example UP means an interface is in service.

Mac Address The Machine Access Control (MAC) address of the Network Interface Card. This attribute is empty when the agent runs in Solaris local zones.

Maximum Transmission Unit The maximum packet size (in bytes) for the specified network interface. This is a fixed value. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the minimum, maximum or average packet size used by a network interface. This information can help you determine the size used by a network interface.

This information can help you determine the data transfer capabilities of various network interfaces, and alleviate bottlenecks by re-routing traffic from devices that are displayed to be overloaded, to other network interfaces that might be able to handle additional data traffic.

Network Interface Name Identifies the network interface adapter. `Aggregate` is a valid value. Valid entries are simple text string, alphanumeric comprised of "Interface Name, Unit Number" where:

- The name is a two-character representation of the adapter, based on the hardware, operating system, and installation procedure.
- The unit represents the physical adapter number installed in the system with a typical range of 0-7.

Example On an AIX system, typical network adapters are represented as follows:

en	= Ethernet
lo	= Loopback
tr	= Token Ring
sl	= SLIP

Other operating systems might refer to the adapter type in a different manner. For example, on SunOS, ethernet adapters are typically represented by **le**.

On HP-UX, you might see **ethernet** represented as **lan**.

The possible combinations based on hardware, operating systems and installation options are virtually impossible to list. Ask your local System Administrator for assistance in determining the specific adapter types installed on your system.

Output Errors The number of packet transmission errors by the network interface during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Output Packet Errors Percent Portion of packets sent over the previous 30-second interval with errors. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Packet Collisions Percent Portion of packets sent over the previous 30-second interval with collisions. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Example If a low number of packets are being received, data traffic might need to be re-routed.

Packets Received The number of packets received by the interface during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Packets Transmitted The number of packets transmitted by the interface during the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Received Megabytes Second The number of megabytes received per second averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Received Megabytes (Total) The number of megabytes received on the interface (total). Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Total Packets Received The number of packets received since the network interface was initialized. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Total Packets Received (superseded) The number of packets received since the network interface was initialized. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Total Packets Transmitted The number of packets transmitted since the network interface was initialized. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum.

Total Packets Transmitted (superseded) The number of packets transmitted since the network interface was initialized. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Example A high value might indicate an overloaded interface. A low value might indicate a device that is not being used much, which can carry an additional load, if required.

Transmitted Megabytes Second The number of megabytes sent per second averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Transmitted Megabytes (Total) The number of megabytes transmitted on the interface (total). Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Type The type of network interface card. The following values are valid:

Ethernet	X.25	IEEE 802.5 Token-Ring	PPP
Other	Localtalk	SLIP	Not Available
Unknown	Multi-protocol Over Frame Relay	Character Asynchronous Protocol	ATM Classical IP Interface
Infiniband	HDH Arpanet	ISO 802.6 MAN	Hyperchannel
SMDS	SONET	802.2 LLC	Frame Relay
Federation Switch	SP Switch	Virtual Interface	DS3/T3
ULTRA	ISO 8802/3 and Ethernet	ISO HDLC Protocol	HIPPI
Frame Relay LAPF	100-Based VG Token Ring	IEEE 802.4 Token-Bus	MODEM
Fibre Channel Interface	100-Based VG Ethernet	PTP Serial	IEEE 802.3 Network
Proprietary Virtual	IBM Channel to Channel Adapter	HSSI	T1/E1
Fiber Distributed Data Interface	Character Synchronous Protocol	ISDN	XNS
ATM	Proprietary Multiplexing	100 Base-T	V.35
SDLC	X.25 LAPB	Not Collected	Software Loopback

NFS / RPC Statistics attributes

The NFS / RPC Statistics group attributes refer to Network File System and Remote Procedure Call call rates and errors. These attributes fall into four subgroups:

- NFS Client attributes report on calls from the managed system to NFS servers.
- NFS Server attributes report on NFS calls to the managed system. The agent reports these calls only when the managed system is an NFS server.
- RPC Client attributes report on calls from the managed system to RPC servers.
- RPC Server attributes report on PRC calls to the managed system. The agent reports these calls only when the managed system is an RPC server.

NFS Client Calls The number of calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to show the amount of NFS traffic. If the value is high, it might mean a client is flooded with call requests.

NFS Client Calls Rejected The number of calls rejected by a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute in conjunction with the NFS Client Calls attribute to determine the proportion of calls rejected by the NFS server.

NFS Client File Creates The number of File Creates calls made to a server during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client File System Statistics Calls The number of file statistics calls made within the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Get Attribute Calls The number of calls made to determine what type of file is being called. For example, a text file or an executable file. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Link Calls The number of hard link reports made by a server during a pre-defined time interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Lookups The number of Lookups requests made by an NFS server during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and

-2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the number of Make Directory requests handled by an NFS server during the monitoring interval.

NFS Client Make Directory Calls The number of calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Null Calls The number of calls generated for checking connectivity to a server. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Read Calls The number of Read Directory calls read by a server during a monitored interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine how many call requests to read a file were received by an NFS server over a period of time.

NFS Client Read Directory Calls The number of Read Directory calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Read Link Calls The number of calls received by an NFS server to read a linked file during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine how many link-call requests an NFS server received over a period of time.

NFS Client Rejected Call Percentage The percentage of NFS calls rejected by a server during a monitoring interval. Valid entry is an integer in the range 0 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Use this attribute to monitor unacceptable rates of NFS call rejection and to determine whether the server or a particular client is causing network problems. If the server is experiencing problems, all calls are rejected. If a client is experiencing problems, it alone has its calls rejected. The rejection percentage might increase dramatically after reaching a critical threshold, since rejections require re-transmission.

NFS Client Remove Directory Calls The number of Remove Directory calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Remove File Calls The number of Remove File calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and

-2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Rename File Calls The number of Rename File calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Root Calls The number of NFS calls made to the server by the root during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the number of calls made by the root (Super-User) account versus calls made by all users.

NFS Client Set Attribute Calls The number of NFS calls made to set the attributes of a file during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Symbolic Link Calls The number of Symbolic Link calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Write Cache Calls The number of Write Cache calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Client Writes The number of Writes calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine how many write requests an NFS server received over a pre-specified period of time.

NFS Server Calls The number of Server Calls made from an NFS server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to show the amount of NFS traffic. If the value is high, it might mean a server is flooded with call requests.

NFS Server Calls Rejected The number of calls rejected by a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648

indicates Value_Exceeds_Minimum. Use this attribute in conjunction with the NFS Server Calls attribute to determine the proportion of calls rejected by the NFS server.

NFS Server File Creates The number of File Creates calls made to a server during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server File System Statistics Calls The number of file statistics calls made within the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server Get Attribute Calls The number of calls made to determine what type of file is being called during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

NFS Server Get Attr Percent The portion of calls made over the previous 30-second interval to determine what type of file is being called. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

NFS Server Link Calls The number of hard link reports made by a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server Lookups The number of Lookups requests made by an NFS server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the number of Lookups requests handled by an NFS server during the monitoring interval.

NFS Server Make Directory Calls The number of Make Directory calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server Null Calls The number of calls generated for checking connectivity to a server. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server Read Calls The number of Read Directory calls read by a server during a monitored interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648

indicates Value_Exceeds_Minimum. Use this attribute to determine how many call requests to read a file were received by an NFS server over a period of time.

NFS Server Read Directory Calls The number of Read Directory calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server Read Link Calls The number of calls received by an NFS server to read a linked file during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine how many link-call requests an NFS server received over a period of time.

NFS Server Read Link Percent The portion of lookup requests made by an NFS server over the previous 30-second interval that read a linked file. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

NFS Server Read Percent The portion of client calls made to a server over the previous 30-second interval that were reads. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

NFS Server Rejected Call Percentage The percentage of NFS calls rejected by a server during a monitoring interval. Valid entry is an integer in the range 0 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Use this attribute to monitor unacceptable rates of NFS call rejection and to determine whether the server or a particular client is causing network problems. If the server is experiencing problems, all calls are rejected. If a client is experiencing problems, it alone has its calls rejected. The rejection percentage might increase dramatically after reaching a critical threshold, since rejections require re-transmission.

NFS Server Remove Directory Calls The number of Remove Directory calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server Remove File Calls The number of Remove File calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server Rename File Calls The number of Rename File calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server Root Calls The number of NFS calls made to server root during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the number of calls made by the root (Super-User) account versus calls made by all users.

NFS Server Set Attribute Calls The number of NFS calls made to set the attributes of a file during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server Symbolic Link Calls The number of Symbolic Link calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server Write Cache Calls The number of Write Cache calls made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute with other Call-type attributes for analyzing NFS traffic and correcting problems on your network.

NFS Server Write Percent The portion of client calls made to a server over the previous 30-second interval that were writes. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

NFS Server Writes The number of Writes made to a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine how many write requests an NFS server received over a pre-specified period of time.

NFS Version The version of NFS from which metrics were obtained, or Aggregate if metrics from more than one version are included. Only Aggregate for the NFS attributes is supported. Note: the value -1 indicates Not Available, -2 indicates Not Collected, and 0 indicates Aggregate. Other values are v1, v2, v3, and v4.

RPC Client Bad Xid Replies Percent The portion of RPC server calls over the previous 30-second interval that did not match the call. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

RPC Client Bad Xid Replies Limited Percent The portion of rejected RPC server calls over the previous 30-second interval that did not match the call. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

RPC Client Calls The number of RPC client calls per second averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

RPC Client Calls Rejected by Server The number of calls made by a client to a server that the server rejected, during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute in conjunction with the Calls attribute to determine the proportion of calls rejected by the RPC server. If the value is high, it might mean there is excessive noise on the network, which causes bad datagrams to occur, or a server might be flooded with call requests.

RPC Client Calls Rejected by Server Percent The portion of calls made by a client to a server over the previous 30-second interval that the server rejected. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

RPC Client Calls Rejected Percent The portion of RPC client calls over the previous 30-second interval that were rejected by the server. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

RPC Client Calls Retransmitted The number of RPC packets retransmitted to an RPC server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine if a large number of calls are being retransmitted. If so, your server might be overworked or there might be a network problem.

RPC Client Calls Retransmitted Limit Percent The portion of timed-out RPC client calls over the previous 30-second interval that were retransmitted. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

RPC Client Calls Retransmitted Percent The portion of RPC packets over the previous 30-second interval retransmitted to an RPC server. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

RPC Client Calls Timed Out The number of times an RPC call from the managed system timed out before the RPC server replied. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine if the server is failing to acknowledge calls received. If the server is overworked, you might need to re-route network traffic.

RPC Client Calls Timed Out Percent The portion of RPC client calls over the previous 30-second interval that timed out. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

RPC Client Replies Not Matching Calls The number of times the managed system received replies from an RPC server that did not match calls. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine if a reply or acknowledgment from the server matches a request made by a client. A failure to match the request might mean there is noise on the network.

RPC Client Times Authentication Refreshed The number of times the managed system had to resend the authentication information for an RPC call during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648

indicates Value_Exceeds_Minimum. Use this attribute to count the number of times an authorization is refreshed. This attribute helps you verify client authorization for making a request by periodically requesting an electronic handshake from the client.

RPC Client Times Call Wait On Busy The number of times the initial bind for an RPC call from the managed system had to wait because of a busy server. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to show the amount of NFS traffic. When a client sends a call request to a server, it gives the server a certain amount of time to respond before re-sending the call. The amount of time varies from system to system. If the Times Call Wait On Busy value is high, it might indicate that the server is overworked. You might want to re-route call requests to another server.

RPC Server Calls The number of RPC server calls per second averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

RPC Server Calls Rejected The number of RPC calls from the managed system that were rejected by a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine if the server is rejecting a large number of calls. If only a few calls are being rejected, it might be a client-specific problem. If many calls are being rejected, it might be a problem with your server.

RPC Server Calls Rejected Percent The portion of RPC server calls over the previous 30-second interval that were rejected by the server. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

RPC Server Dup Checks The number of RPC server calls per second serviced from the duplicate request cache averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

RPC Server Dup Reqs The number of duplicate RPC server calls per second averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

RPC Server Dup Reqs Percent The portion of RPC server calls over the previous 30-second interval that were duplicate requests. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

RPC Server Packets Too Short The number of incomplete RPC packets that were too short in length received by a server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine if the server is having problems processing packet data. If the packet size does not match the size stated in the packet header, there might be noise on the system.

RPC Server Packets with Malformed Header The number of RPC packets that had malformed headers and were received by the server during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine if there is noise on the system. The server cannot validate a packet or where it came from if, due to a malformed header, it cannot acknowledge the sender. This decreases the efficiency of the network. Try checking server connections. Another cause might include extraneous network noise.

RPC Server Times RPC Packet Unavailable The number of times a server attempted to receive a packet when none was available, during a monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System Name The managed system name. The form should be *hostname:agent_code*.

Examples include spark:KUX or deux.raleigh.ibm.com:KUX.

In workspace queries, this attribute should be set equal to the value $NODE$ in order to populate the workspace with data. This attribute is generally not included in situations, unless there is a need to customize the situation for a specific managed system.

Timestamp The date and time the agent collects information as set on the monitored system. The timestamp for SCAN and STR functions is in the CYYMMDDHHMMSSmmm format (as in 1020315064501000 for 03/15/02 06:45:01) where:

C	= Century (0=20th)
YY	= Year
MM	= Month of the Year (01-12)
DD	= Day of the Month (01-31)
HH	= Hour, in 24-hour time (00-23)
MM	= Minute
SS	= Second
mmm	= Millisecond

Ping attributes

The Ping attributes refer to target characteristics such as target host and ping result. Historical information is available for the Ping table for users interested in trending server response times. However, to enable history collection for this attribute group, a list of monitored (pinged) servers must be specified. The list is specified through an environment variable - "KUX_PINGHOSTLIST" in the `ux.ini` file in the IBM Tivoli Monitoring config directory. For example:

```
KUX_PINGHOSTLIST='/opt/IBM/ITM/config/kuxpinghosts'

sample content of kuxpinghosts:
#
# hosts pinged for availability from this agent
```

```
#
server1.domain.com
server2
server4
```

Once the kuxpinghosts file has been created on the monitored system and the parm entry has been made in the ux.ini file, you can create a table to display the UNIX Ping attributes. The returned rows would be for each of the hosts specified in the kuxpinghosts file on the monitored system.

You can also create a situation using the UNIX Ping attributes. If the situation does not contain a value in the Target Host attribute field, then the systems identified in the kuxpinghosts file is pinged. If the situation contains a host name or ip address in the Target Host attribute field in the situation, then this Target Host name overrides the use of the kuxpinghosts file (and the hosts specified therein).

To create this situation, in the situation editor complete the following steps:

- Create the situation.
- Set the System Node to $NODE$.
- Set the Target_Host or leave it blank to use a ping hosts file.
- Select **Advanced** and the select **Display Item**.
- Set the Display Item to **Target_Host**.

This creates a situation that generates an event and report the target host's IP in the situation. It is especially useful if you set up multiple Ping queries.

Note: This agent does not support the use of "Value of expression IN" for Ping Attributes.Situations can be successfully defined using the *IN method, but the situation will only evaluate against the first entry in the list.

Ping Result Result from pinging the host. The following values are valid: Successful, Unsuccessful, and Not Collected and Not Available.

Server Response Time Ping Operation response time in milliseconds. The following values are valid: a numerical value in milliseconds, Not Collected and Not Available.

System Name The managed system name.

Target Host Host name or IP address of the host to be pinged.

Timestamp The date and time the agent collects information as set on the monitored system.

Print Queue attributes

The Print Queue attributes refer to print queue characteristics.

Device Name The name of a device associated with this queue.

Print Queue Depth The number of jobs in the print queue. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Print Queue Name The name of the print queue.

Print Queue Job Size The number of kilobytes in the print queue, including copies. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum

Print Queue Status The status of the print queue.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Process attributes

The Process Attributes Group refers to process operations such as command executed, CPU utilization, real memory usage, and execution state.

Note: Do not collect history from this attribute group due to the volume of information that might be sent to the Tivoli Data Warehouse.

Special information about the Flag field (Solaris only)

The Flag field of the UNIX Process report contains hexadecimal and additive flags. These flags are available for historical purposes only, and contain no information regarding the current status of your monitored process. These fields are not relevant on Solaris systems. For additional information about the Flag field, see the man pages for your operating system.

Base Command The command that initiated a process.

Child System CPU Time The time spent in system and user mode by the child of this process. Valid entry is a numeric time string with a format of DDDdHH:MM:SS where:

DDD	= Days to a maximum of 999
HH	= Hours
MM	= Minute
SS	= Second

Example:To express 45 days, 1 hour, 5 minutes, and 30 seconds, enter 045d01:05:30. Note: Not Available indicates a value of 0.

Use this attribute to determine which processes have children consuming a large amount of CPU time and take corrective action.

Child User CPU Time The time spent in user mode by the child of this process. Valid entry is a numeric time string with a format of DDDdHH:MM:SS where:

DDD	= Days to a maximum of 999
HH	= Hours
MM	= Minute
SS	= Second

Example: To express 45 days, 1 hour, 5 minutes, and 30 seconds, enter 045d01:05:30. Note: Not Available indicates a value of 0.

Command The command that initiated a process. Valid entry is a simple text string, alphanumeric with a maximum 32 characters. Use this attribute to determine which command initiated a process.

Command (Unicode) The command that initiated a process. Use this attribute to determine which command initiated a process.

Example: An example of a command that initiates a process would be run.

Context Switch The number of CPU context switches for this process. A context switch occurred when a process voluntarily giving up the processor before its time slice was completed. This usually occurs while the process waits for a resource. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum. Use this attribute to monitor for context switches. Excessive context switches might indicate too many waits for resources. Available on Solaris and AIX only.

Context Switch (Superseded) The number of CPU context switches for this process. A context switch occurred when a process voluntarily giving up the processor before its time slice was completed. This usually occurs while the process waits for a resource. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to monitor for context switches. Excessive context switches might indicate too many waits for resources. Available on Solaris and AIX only.

Count of Processes The count of processes with the same name. On systems with AIX WPARs or Solaris Zones where all processes from the containers are visible, the process count includes all processes from all virtual containers and the count is not respective to each WPAR or Zone. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum.

CPU ID The ID of the processor on which the process is running. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the processor on which a process is running. Available on Solaris and HP-UX only.

CPU Pct The percentage of CPU used by this process. As this value is normalized, it does not exceed 100 percent, even if more than one processor is installed. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Use this attribute to determine which processes are using the most CPU time. High CPU percent might indicate a runaway or long running process.

For example, enter 50 to represent 50.00%, or 50.34 to represent 50.34%.

CPU Time The time, in seconds, the CPU has been utilized. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

CPU Utilization The numerical value indicating the relative CPU intensity of a process. The CPU Utilization attribute represents the number of times a process uses the CPU over a period of 20 system clock ticks. The system decays this value after each 20 clock-tick period by dividing the number by 2. The system uses CPU Utilization to determine process priority. Large values indicate a CPU intensive process and result in lower process priority. Small values indicate an I/O intensive process and result in a more favorable priority. Valid entry is a numeric value in the range 0 to 999. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Use this attribute to check process if you suspect it is using the CPU so much that the CPU is not available to anything else. This can cause network response time to be sluggish.

Example: A high value indicates a CPU-intensive process. A low value indicates an I/O-intensive process.

Effective Group ID The effective GROUP ID. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the effective group ID for this process. Available on all platforms.

Effective Group Name The effective group name of the process.

Effective User ID The effective USER ID. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the effective user ID for this process. Available on all platforms

Effective User Name The effective user name of the process.

Elapsed Time The elapsed time for the process. Note: the value -0 indicates Not Available. Valid entry is a numeric time string with a format of DDDdHH:MM:SS where:

DDD	= Days to a maximum of 999
HH	= Hours
MM	= Minute
SS	= Second

Use this attribute to determine how long this process has been running.

Example: To express 45 days, 1 hour, 5 minutes, and 30 seconds, enter **045d01:05:30**

Entry Address The virtual memory address of a process. This address corresponds to the segment number of the process stack. Valid entry is a hexadecimal string. Check with you local System Administrator for information on how to use this attribute. **Note:** On 64-bit systems, only the low-order part of the address is used.

Example: The virtual memory address of a process varies from process to process.

Event Waited On The memory address of an event (if any) on which a process is waiting. A process must have this information before it can run. Valid entry is a

simple text string or hexadecimal value depending on the operating system with a maximum string length of 8. This information is specific to your particular network.

Example: On AIX, the word EVENT is an example of what is displayed for this attribute.

Execution State The execution state of a process. For valid entries, use one of the following codes to indicate an execution state:

0	= Non-existent
A	= Active
I	= Intermediate
O	= Running
R	= Runable
S	= Sleeping
T	= Stopped
W	= Waiting
X	=Growing
Z	= Zombie

Use this process to determine the state of a particular process. If a process is waiting, there might be an excessive amount of network traffic, or a process might be taking a long time to complete. Further investigation might be needed. Check with your local system administrator to determine what corrective action to take.

Flag The hexadecimal value associated with a process. Valid entry is a hexadecimal value with a maximum string length of 8. The meaning of a flag depends on the type of UNIX system you are monitoring. The Flag field of the UNIX Process report contains hexadecimal and additive flags. These flags are available for historical purposes only, and contain no information regarding the current status of your monitored process. These fields are not relevant on Solaris systems. For additional information about the Flag field, please refer to the man pages for your operating system.

Group Name The group name of the process owner.

Heap Size The size of the heap for this process expressed in bytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the heap size for a process. Excessive heap size might indicate a memory leak. Available on Solaris only.

Involuntary Context Switch The number of involuntary context switches for the process. An involuntary context switch occurs when a higher priority process ran or because the current process exceeded its time slice. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum. Use this attribute to monitor for involuntary context switches. Excessive involuntary context switches might indicate function problems in a process. Available on Solaris and AIX only.

Involuntary Context Switch (Superseded) The number of involuntary context switches for the process. An involuntary context switch occurs when a higher priority process ran or because the current process exceeded its time slice. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to monitor for involuntary context switches. Excessive involuntary context switches might indicate function problems in a process. Available on Solaris and AIX only.

Major Fault The number of major faults requested by this process. A major fault requires disk access. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum. Use this attribute to monitor for major faults. Excessive major faults might indicate memory shortage.

Major Fault (Superseded) The number of major faults requested by this process. A major fault requires disk access. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to monitor for major faults. Excessive major faults might indicate memory shortage.

Mem Pct The percentage of system memory used by this process. Valid entry is a numeric value in the range 0 to 100.00 to two decimal places. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Use this attribute to monitor memory usage by a process. Processes with high memory percent leads to memory shortage and cause system performance problems.

Example: Enter 50 to represent 50.00%, or 50.34 to represent 50.34%.

Minor Fault The number of page reclaims for the process. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum. Use this attribute to determine the number of minor faults occurred in any processes.

Minor Fault (Superseded) The number of page reclaims for the process. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the number of minor faults occurred in any processes.

Nice Value The requested execution priority of a process, in relation to other processes. The higher the nice value, the lower the priority of the command. The nice value, plus the minimum user process priority level equals the priority of the process. The range of nice values varies among UNIX systems. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. The range of nice values varies among UNIX systems. Check with your local system administrator for information concerning the range of nice values for your system.

Page Space Used (AIX) The amount of page space used by the process private data (4K pages). Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Parent Process ID The unique numerical identifier of a process. The process that invoked the forked system call is the parent process, and the newly created process is the child process. Every process has one parent process, but a process can have several children. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the PPID for this process.

Priority The current execution priority value. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. The priority equals the nice value of the process plus the minimum priority value assigned to all user processes. The higher the priority value, the lower the priority of the command.

Example: A value of 245 indicates the process is scheduled to be the 245th process ran. This is usually considered to be a low scheduling priority.

Process Command A command string including the arguments up to 100 characters in length. Valid entry is a simple text sting with a maximum 100 characters. Use this attribute to determine which command initiated this process.

Process Command (Unicode) A command string including the arguments up to 768 bytes in length. Use this attribute to determine which command initiated this process.

Process Filter (Unicode) A regular expression to be applied to the Process Command (Unicode) attribute. The maximum allowable length is 256 characters. The following values are valid:

- Java_processes_(.*java.*)
- IBM_Java_processes_entry_method_only_(.*java.*(com.ibm.*))
- System_Admin_installed_processes_(/usr.*)

Process Group Leader ID The process group leader PID. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the process group leader ID for this process. Available on all platforms

Process ID The numerical process ID assigned to a process. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the process ID for this process. Process ID values vary from system to system.

Read/Write The number of characters read and write by this process. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 9223372036854775807 indicates Value_Exceeds_Maximum, and -9223372036854775808 indicates Value_Exceeds_Minimum. Use this attribute to determine the number of read and write completed by this process.

Read/Write (Superseded) The number of characters read and write by this process. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647

indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the number of read and write completed by this process.

Real Group ID The real group ID for this process. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the real group id for this process. Not available on HP-UX.

Resident Data Size (AIX) The amount of resident physical memory used by the process private data (4K pages). Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Resident Text Size (AIX) The amount of resident physical memory used by the process code (4K pages). Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Scheduling Class The scheduling class for this process. Valid entry is a simple text sting with a maximum 8 characters. Use this attribute to determine the scheduling class of this process. Not available on AIX and HP-UX.

Session ID The real session ID for this process. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Size (KBytes) The resident set size of the process, in kilobytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine which processes are using too much memory. Excessive resident set size might lead to memory shortage and cause system performance problems.

Stack Size The size of the stack for this process. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine which processes are using too much stack size.

Start Time The time when the process was started. Note: the value -0 indicates Not Available. Valid entries are in the format CYYMMDDHHMMSSmmm, where:

C	= Century (0=20th)
YY	= Year
MM	= Month of the Year (01-12)
DD	= Day of the Month (01-31)
HH	= Hour, in 24-hour time (00-23)
MM	= Minute
SS	= Second
mmm	= Millisecond

Example: A value of 0951009130500000 indicates the agent collected the data on October 9, 1995 at 1:05 p.m.

System CPU Time The system time spent executing this process. Note: the value -0 indicates Not Available. Valid entry is a numeric time string with a format of DDDdHH:MM:SS where:

DDD	= Days to a maximum of 999
HH	= Hours
MM	= Minute
SS	= Second

Use this attribute to monitor the system CPU time spent by any processes. Excessive system CPU time might indicate a runaway or long running process.

Example: To express 45 days, 1 hour, 5 minutes, and 30 seconds, enter **045d01:05:30**

System Name The managed system name.

Terminal Device The name of the terminal device that started a process. Valid entry is a simple text string with a maximum 8 characters. Terminal names vary from system to system. Check with your local system administrator for a complete list of all terminals in your system.

Thread Count The total number of threads for the process. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the number of threads in this process for information only. Not available on HP-UX.

Time The total amount of CPU time that a process has consumed. Should this value become large, it might indicate a runaway or long-running process. Valid entry is a numeric time string with a format of MMMMM:SS where:

MMMMM	= Minute
SS	= Second

Example: To express 1 hour, 5 minutes, and 30 seconds, enter 00065:30

Use this attribute to identify runaway or long-running processes.

Timestamp The date and time the agent collects information as set on the monitored system.

Total Child CPU Time The sum of the child CPU time (user + system) spent executing the process. Valid entry is a numeric time string with a format of DDDdHH:MM:SS where:

DDD	= Days to a maximum of 999
HH	= Hours
MM	= Minute
SS	= Second

Example: To express 45 days, 1 hour, 5 minutes, and 30 seconds, enter **045d01:05:30**. Note: Not Available indicates a value of 0.

Total CPU Percent The percentage of CPU used since the process was started. As this value is normalized, it does not exceed 100 percent, even if more than one processor is installed. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Use this attribute to identify which processes are using the most CPU time. Excessive total CPU percent might indicate a runaway or long running process.

Example: Enter 50 to represent 50.00%, or 50.34 to represent 50.34%

Total CPU Time The total CPU time (user + system) spent on the process. Note: the value -0 indicates Not Available. Valid entry is a numeric time string with a format of DDDdHH:MM:SS where:

DDD	= Days to a maximum of 999
HH	= Hours
MM	= Minute
SS	= Second

Example: To express 45 days, 1 hour, 5 minutes, and 30 seconds, enter 045d01:05:30

Use this attribute to identify which processes are using the most CPU time. Excessive CPU time might indicate a runaway or a long running process.

Type The type of UNIX operating system residing on a monitored host.

User CPU Time The user CPU time spent executing the process. Note: the value -0 indicates Not Available. Valid entry is a numeric time string with a format of DDDdHH:MM:SS where:

DDD	= Days to a maximum of 999
HH	= Hours
MM	= Minute
SS	= Second

Example: To express 45 days, 1 hour, 5 minutes, and 30 seconds, enter **045d01:05:30**

Use this attribute to identify which processes are using the most CPU time. Excessive CPU time might indicate a runaway or a long running process.

User ID The numerical user ID of the owner of a process. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to identify the owner of a process.

Example: The numeric identification number varies from system to system and user to user. An example of a user ID could be 48765.

User Name The login name of the user based on UID. Valid entry is a simple text sting with a maximum 32 characters. Use this attribute to identify the owner of a process.

User Name (Unicode) The login name of the user based on UID. Use this attribute to identify the owner of a process.

Virtual Size The size of the virtual memory used by this process, in kilobytes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the size of the virtual memory used by a process. Excessive virtual memory size might indicate a memory leak.

Wait CPU Time The time spent waiting for the CPU. Valid entry is a numeric time string with a format of DDDdHH:MM:SS where:

DDD	= Days to a maximum of 999
HH	= Hours
MM	= Minute
SS	= Second

Example: To express 45 days, 1 hour, 5 minutes, and 30 seconds, enter **045d01:05:30**. Note: Not Available indicates a value of 0.

Use this attribute to determine the time spent waiting for CPU. Excessive wait for CPU time might indicate a system problem. Available on Solaris only.

Wait Lock Time The time spent waiting for locks to release. Valid entry is a numeric time string with a format of DDDdHH:MM:SS where:

DDD	= Days to a maximum of 999
HH	= Hours
MM	= Minute
SS	= Second

Example: To express 45 days, 1 hour, 5 minutes, and 30 seconds, enter 045d01:05:30. Note: Not Available indicates a value of 0.

Use this attribute to determine the time spent waiting for a lock. Excessive wait for lock time might indicate a resource concurrence problem. Available on Solaris only.

WLM Name (AIX) The WLM class name to which the process belongs. Valid entry is a simple text string, alphanumeric with a maximum 100 characters. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

WPAR Name (AIX) The name of the WPAR. Valid entry is a simple text string, alphanumeric with a maximum 100 characters. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Zone ID (Solaris) The ID of the Solaris zone. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Zone Name (Solaris) The name of the Solaris zone. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

SMP CPU attributes

The SMP CPU Group attributes refer to multi-processor characteristics such as cross-calls, thread migrations, and system calls.

Avg CPU Busy 1 The average CPU busy time over the last one minute. Displays N/C (not collected) if the agent has been up for less than one minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg CPU Busy 5 The average CPU busy time over the last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg CPU Busy 15 The average CPU busy time over the last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg CPU Busy 60 The average CPU busy time over the last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg CPU System 1 The average system CPU time over the last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg CPU System 5 The average system CPU time over the last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg CPU System 15 The average system CPU time over the last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg CPU System 60 The average system CPU time over the last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg CPU User 1 The average user CPU time over the last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg CPU User 5 The average user CPU time over the last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the

value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg CPU User 15 The average user CPU time over the last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg CPU User 60 The average user CPU time over the last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Context Switches The CPU context switches rate per second during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use to determine system workload or per processor workload of the SMP system. Available on Solaris only.

Context Switches per Sec (AIX) The process context switches on this processor per second during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use to determine system workload or per processor workload of the SMP system.

CPU Busy (Percent) The sum of the System CPU and User CPU attributes in percent. Valid entries in the range 0 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Use to determine system workload or per processor workload of the SMP system.

CPU ID The processor ID. Use this attribute to determine the processor ID. In an SMP system with more than one processor, the CPU report shows CPU ID as “aggregate” on the first row. This means the data row return aggregated CPU statistics.

CPU Status The current status of the processor. Valid entries include 0 = off-line and 1 = on-line. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

CPU Time The time the CPU has been utilized. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

CPU Usage The sum of the percent user and percent system time of the CPU averaged over the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Cross Calls Inter-processor cross-calls rate per second during the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the inter-processor cross reference call rate of the system or per processor of the SMP system. Available on Solaris only.

Idle CPU (Percent) Percentage of idle CPU time during the sampling period. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Use this attribute to determine how efficiently the entire system or each processor of the SMP system is operating. The Idle CPU value must be low if the system load is heavy, and high if the system load is light. If the system load is heavy and the Idle CPU value is high, an I/O problem might exist. If the Idle CPU value is small, or zero, and the User percentage is larger (greater than 30%), the system might be compute-bound or in a loop. Available on all platforms.

Interrupts Interrupts rate per second over the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the devices interrupts rate of the system or of each processor of the SMP system. Available on Solaris only.

Interrupts As Threads Interrupts as thread (not counting interrupts) rate per second during the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the rate of interrupts as threads (below block) of the system or of each processor of the SMP system. Available on Solaris only.

Involuntary Context Switches Involuntary context switches rate per second during the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Logical Context Switches (AIX) The number of logical context switches per second during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use to determine system workload or per processor workload of the SMP system.

Major Faults Major faults rate per second during the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the rate of page faults that need disk access of the system or of each processor of the SMP system. Available on Solaris only.

Minor Faults Minor faults rate per second during the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the rate of page faults where the pages faulted is located in memory, usually on the inactive list for the entire system or for each processor of the SMP system. This is available on Solaris only.

Physical Consumption (AIX) The number of physical CPU units consumed by this logical CPU. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Spins On Mutexes Spins on mutexes (locks not acquired on try) rate per second during the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648

indicates Value_Exceeds_Minimum. Use this attribute to determine the spins on mutexes rate of the system or of each processor of the SMP system. Available on Solaris only.

Spins On RW Locks Spins on read/write locks (locks not acquired on first try) rate per second during the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the spins on read write locks rate of the system or of each processor of the SMP system. Available on Solaris only.

System Calls System calls rate per second during the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the system calls rate of the system or of each processor of the SMP system. Available on Solaris only.

System CPU (Percent) Percent of system CPU time during the sampling period. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Use this attribute to determine the percentage of system or per processor CPU time devoted to executing UNIX system kernel code. System CPU time includes time spent executing system calls and performing administrative functions. Available on all platforms.

System Name The managed system name.

Thread Migrations Thread migrations to another processor rate per second during the sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum. Use this attribute to determine the rate of thread migrations to another processor of the system or of each processor of the SMP system. Available on Solaris only.

Timestamp The date and time the agent collects information as set on the monitored system.

User CPU (Percent) Percent of user CPU time during the sampling period. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Use this attribute to determine the percentage of system or per processor CPU time devoted to user processes. User CPU time includes time spent executing both user program and library functions. It does not include CPU time spent executing system calls. The ratio between user and system CPU time varies, depending on the kinds of programs executing. If user CPU is extremely high and adversely affecting system performance, you might want to determine which user programs are preventing the CPU from functioning at its normal speed. Available on all platforms.

Wait I/O Percent of wait I/O CPU time during the sampling period. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected. Use Wait I/O to indicate how effectively the system or a processor is using disks. Available on all platform.

Solaris Zones attributes

The Solaris Zones attributes refer to zone characteristics such as init process ID, path, and zone ID.

Note: The prctl command allows you to get or to set the resource controls of running processes, tasks, and projects. On Solaris local zones, the Monitoring Agent for UNIX OS kuxagent process can result in a local zone crash because of a Solaris bug affecting the prctl command. To avoid this issue, set the environment variable KUX_PRCTL_OFF=TRUE in the `ux.ini` Agent configuration file to disable the collection of CPUSHARES and SHAREPCT data for Solaris zones.

Capped CPU The cap configured in the zone on CPU units. Note: the value -1 indicates Not Available and the value -2 indicates Uncapped.

Capped Memory The cap configured in the zone on physical memory (KB). Note: the value -1 indicates Not Available and the value -2 indicates Uncapped.

CPU Share PCT The percent of the processor set available to this zone. Note: the value -1 indicates Not Available and the value -2 indicates Not Collected.

CPU Shares The weight used by the fair share scheduler to control CPU usage. Note: the value -1 indicates Not Available, -2 indicates Not Collected, and 2147483647 indicates Value_Exceeds_Maximum.

Dedicated CPU The number of CPUs that are assigned for zone exclusive use. Note: the value -1 indicates Not Available and the value -2 indicates None.

Init Process ID The process ID of the init process for this zone. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Name The name of a zone. Valid entry is a simple text string, alphanumeric with a maximum length 192 characters. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Path The path to the root directory of the zone.

Physical Memory The physical memory in kilobytes used by all processes in the zone. Note: the value -1 indicates Not Available, -2 indicates Not Collected, and 2147483647 indicates Value_Exceeds_Maximum.

Pool ID The numeric pool ID this zone is bound to. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Scheduler The scheduler used by the zone. The following values are valid: Fair_Share_Scheduler, Time_Share, Interactive, Fixed_Priority, System, and Real_Time.

Status The status of a zone. Note: the value -1 indicates Not Available and -2 indicates Not Collected. The following values are valid:

- Uninitialized
- Ready

- Booting
- Running
- Shutting Down
- Empty
- Down
- Dying
- Dead
- Initialized

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Total CPUs The number of CPUs that are in the processor set. Note: the value -1 indicates Not Available, -2 indicates Not Collected, and 2147483647 indicates Value_Exceeds_Maximum.

Virtual Memory The virtual memory in kilobytes used by all processes in the zone. Note: the value -1 indicates Not Available, -2 indicates Not Collected, and 2147483647 indicates Value_Exceeds_Maximum.

Zone CPU Usage The CPU usage of all processes in the zone. Note: the value -1 indicates Not Available and the value -2 indicates Not Collected.

Zone ID The full name of a user. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

SP2 System attributes

SP2 System is a virtual attribute group that aggregates data at the Tivoli Enterprise Monitoring Server for the attribute groups of all your instances of the Monitoring Agent for UNIX OS. Any workspace or situation that queries against this attribute group are fulfilled at the Tivoli Enterprise Monitoring Server. This attribute group is useful for monitoring a large number of systems to reduce query time and network load. SP2 System attributes refer to system characteristics such as the amount of available virtual memory, idle CPU percentage, the number of non-block device reads, and load averages.

Note: This attribute group is not historically collected.

Active Virtual Memory (KBytes) The amount of real memory and secondary storage, in kilobytes, currently in use by the system for paging, system users, and caching. On HP-UX and Solaris, this value does not include in-use real memory. This is a deprecated attribute. New queries should utilize the UNIX Memory attributes for clearer definitions. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Ins 1 Minute The average rate of page ins over the last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note:

the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Ins 5 Minute The average rate of page ins over the last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Ins 15 Minutes The average rate of page ins over the last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Ins 60 Minutes The average rate of page ins over the last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Outs 1 Minute The average rate of page outs over the last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Outs 5 Minutes The average rate of page outs over the last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Outs 15 Minutes The average rate of page outs over the last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Outs 60 Minutes The average rate of page outs over the last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Scans 1 Minute The average rate of page scans over the last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Scans 5 Minutes The average rate of page scans over the last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes.

Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Scans 15 Minutes The average rate of page scans over the last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Scans 60 Minutes The average rate of page scans over the last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Processes RunQueue 60 Minutes The average number of processes in the run queue over the last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Block Reads The number of physical block reads over a specified sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Block Writes The number of physical block writes (synchronous+ asynchronous) over a specified sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Boot Time The system boot time on the monitored system. Valid entries are in the format CYYMMDDHHMMSSmmm, where:

C	= Century (0=20th)
YY	= Year
MM	= Month of the Year (01-12)
DD	= Day of the Month (01-31)
HH	= Hour, in 24-hour time (00-23)
MM	= Minute
SS	= Second
mmm	= Millisecond

CPU Busy (Percent) The sum of the System and User CPU attributes in percent. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

CPU Context Switches Number of CPU context switches over the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Device Interrupts The number of non-clock device interrupts over the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Execs Executed The number of execs executed over a specified sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Forks Executed The number of forks executed over the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Free Virtual Memory (KBytes) The amount of unused real memory and secondary storage in kilobytes, currently free for the system to use for paging. On HP-UX and Solaris, this value does not include unused real memory. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Idle CPU (Percent) The percentage of time the CPU is not processing instructions. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

IPv6_Address The IPv6 address corresponding to the system host name. This attribute only supports IPv6 addresses. The following values are valid: Not_Collected and Not_Available among others.

Load Average 1 Min The average number of processes in the run queue of the UNIX kernel during the last minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Load Average 15 Min The average number of processes in the UNIX kernel run queue during the last 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Load Average 5 Min The average number of processes in the UNIX kernel run queue during the last 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Logical Block Reads Number of logical block reads of system buffers during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Logical Block Writes Number of logical block writes of system buffers during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Net Address The Internet Protocol (IP) address of a monitored system, expressed in dotted decimal format. Valid entry is a simple text string, alphanumeric. The following values are valid: Not_Collected and Not_Available among others. This attribute only supports IPv4 addresses.

NonBlock Reads Number of physical block reads (synchronous + asynchronous) during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

NonBlock Writes The number of raw I/O writes over a specified sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Number of System Procs The number of processes running on the system, including both system and user processes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Number of Users Sessions The number of interactive user sessions running. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Faults (Per Sec) The average rate of page faults per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Ins (Per Sec) The average rate per second of page-in requests over a specified sampling period. A page-in request might include multiple pages and gives an indication of the I/O rate on the paging file. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Ins Rate The number of kilobytes that the virtual memory manager pages in per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Outs (Per Sec) The average rate per second of page-out requests over a specified sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Out Rate The number of kilobytes that the virtual memory manager pages out per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Reclaims (Per Sec) The number of times during the monitoring interval that the system removed a page from the queue and used that page for another process. This is the average rate per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Scan Rate (Per Sec) The average rate per second of pages examined over the sampling interval. Virtual memory manager uses a clock-algorithm to implement a pseudo “least recently used (lsu), page replacement scheme. Pages are aged by examination by the clock. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Scanning The number of kilobytes that the virtual memory manager pages scans per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Pages Paged In (Per Sec) The average rate per second of pages that were paged-in from disk to system memory during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Pages Paged Out (Per Sec) The average rate per second of pages paged-out from system memory to disk during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Idle Number of processes currently in idle state. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Runnable Number of processes waiting to be run. (For AIX and HP-UX, this includes both the processes able to be run as well as the running processes.) For AIX, they are the processes in the SACTIVE state: active. For HP-UX, they are the processes in the PS_RUN state: running or waiting for CPU. For Solaris they are the processes in the R (SRUN) state: able to be run, but not currently running. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Running Number of processes currently running on a processor (available for Solaris only). For Solaris, it is the number of processes in the O (SONPROC) state: running. This attribute is not available for AIX and HP-UX because, on these systems, this information is collected as part of the Processes Runnable attribute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Sleeping Number of processes currently in sleep state. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Stopped Number of processes currently in the stopped state. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Threads in Run Queue The total number of processes that can be run (or threads in AIX 4 1 and above) waiting for execution by the CPU. This number does not include processes waiting for I/O or some external event, or processes in a sleeping state. The following values are valid: numeric values in the range 0 to

999. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Threads Waiting The number of processes (or threads in AIX 4 1 and above) waiting for page operations. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Zombie Number of zombie processes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Swap Space Free Amount of swap space (in MB) available. Swap space is usually a disk partition on which page-outs are written. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System Calls Number of system calls made during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System CPU (Percent) The percentage of CPU time devoted to executing UNIX system kernel code. The following values are valid: numeric values expressed as a percentage in the range 1 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

System Name The managed system name.

System Read The number of read and readv system calls during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System Write Number of write and writev system calls over the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Timestamp The date and time the agent collects information as set on the monitored system.

Total Real Memory (KBytes) The total number of kilobytes of physical memory (primary paging memory) on the monitored system. This is a deprecated attribute. New queries should utilize the UNIX Memory attributes for clearer definitions. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Total Virtual Memory (KBytes) The Virtual memory is the total amount memory that is available virtually. The total amount memory that is displayed (virtual) to be available is nothing but the real memory and the part of the disk (paging space) that is available for usage as memory. This is a deprecated attribute. New queries should utilize the UNIX Memory table attributes for more clear definitions. Note:

the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Type The type of UNIX operating system residing on a monitored host. For example, you might see AIX, HPUX, SunOS or OSF1 displayed on your system. The following values are valid:

AIX	The IBM AIX operating system
HPUX	The Hewlett Packard HP-UX operating system
SunOS	Sun Microsystems Solaris I or II operating system

- Maximum length 8 characters

Up Time (Seconds) The number of seconds that a monitored system has been running continuously. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

UpTime The system up time of the monitored system. Valid entries in the format DDDdHH:MM:SS, where:

DDD	= Days
HH	= Hours
MM	= Minutes
SS	= Seconds

User CPU (Percent) The percentage of processor time devoted to user processes. The following values are valid: numeric values expressed as a percentage in the range 0 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Version The version number of a UNIX operating system on the network. Valid entry is a simple text string, alphanumeric with a maximum length of 16 characters.

Virtual Memory Percent Available Percent of virtual memory available. This is calculated by the agent using the formula:

100 - Percent_Virtual_Memory_Used.

The following values are valid: numeric values in the range 0 to100.0 to one decimal place. Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

Virtual Memory Percent Used Percent of virtual memory used. This calculated by the agent using the formula:

Active_Virtual_Memory / Total_Virtual_Memory * 100.

The following values are valid: numeric values in the range 0 to 100.0 to one decimal place. Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

Wait I/O (Percent) The percentage of time the CPU spends waiting for I/O operations. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

Zone ID (Solaris) The ID of the Solaris zone. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Zone Name (Solaris) The name of the Solaris zone. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System attributes

System attributes refer to system characteristics such as the amount of available virtual memory, idle CPU percentage, the number of non-block device reads, and load averages.

Active Virtual Memory (KBytes) The amount of real memory and secondary storage, in kilobytes, currently in use by the system for paging, system uses and caching. On HP-UX and Solaris, this value does not include in-use real memory. This is a deprecated attribute. New queries should utilize the UNIX Memory attributes for clearer definitions. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Ins 1 Minute The average rate of page ins over the last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Ins 5 Minute The average rate of page ins over the last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Ins 15 Minutes The average rate of page ins over the last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Ins 60 Minutes The average rate of page ins over the last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Outs 1 Minute The average rate of page outs over the last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Outs 5 Minutes The average rate of page outs over the last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Outs 15 Minutes The average rate of page outs over the last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Outs 60 Minutes The average rate of page outs over the last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Scans 1 Minute The average rate of page scans over the last 1 minute. Displays N/C (not collected) if the agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Scans 5 Minutes The average rate of page scans over the last 5 minutes. Displays N/C (not collected) if the agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Scans 15 Minutes The average rate of page scans over the last 15 minutes. Displays N/C (not collected) if the agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Page Scans 60 Minutes The average rate of page scans over the last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Avg Processes RunQueue 60 Minutes The average number of processes in the run queue over the last 60 minutes. Displays N/C (not collected) if the agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Block Reads The number of physical block reads over a specified sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Block Writes The number of physical block writes (synchronous+ asynchronous) over a specified sampling period. Note: the value -1 indicates Not Available, -2

indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Boot Time The system boot time on the monitored system. Valid entries are in the format CYYMMDDHHMMSSmmm, where:

C	= Century (0=20th)
YY	= Year
MM	= Month of the Year (01-12)
DD	= Day of the Month (01-31)
HH	= Hour, in 24-hour time (00-23)
MM	= Minute
SS	= Second
mmm	= Millisecond

CPU Busy (Percent) The sum of the System and User CPU attributes in percent. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

CPU Context Switches Number of CPU context switches over the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Device Interrupts The number of non-clock device interrupts over the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Execs Executed The number of execs executed over a specified sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Forks Executed The number of forks executed over the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Free Virtual Memory (KBytes) The amount of unused real memory and secondary storage in kilobytes, currently free for the system to use for paging. On HP-UX and Solaris, this value does not include unused real memory. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Idle CPU (Percent) The percentage of time the CPU is not processing instructions. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

IPv6_Address The IPv6 address corresponding to the system host name. This attribute only supports IPv6 addresses. The following values are valid: Not_Collected and Not_Available among others.

Load Average 1 Min The average number of processes in the run queue of the UNIX kernel during the last 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Load Average 15 Min The average number of processes in the UNIX kernel run queue during the last 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Load Average 5 Min The average number of processes in the UNIX kernel run queue during the last 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Logical Block Reads Number of logical block reads of system buffers during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Logical Block Writes Number of logical block writes of system buffers during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Net Address The Internet Protocol (IP) address of a monitored system, expressed in dotted decimal format. Valid entry is a simple text string, alphanumeric. The following values are valid: Not_Collected and Not_Available among others. This attribute only supports IPv4 addresses.

NonBlock Reads Number of physical block reads (synchronous + asynchronous) during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

NonBlock Writes The number of raw I/O writes over a specified sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Number of CPUs (AIX) The number of logical CPUs that are active. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Number of System Procs The number of processes running on the system, including both system and user processes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Number of Users Sessions The number of interactive user sessions running. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Faults (Per Sec) The average rate of page faults per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Ins (Per Sec) The average rate per second of page-in requests over a specified sampling period. A page-in request might include multiple pages and gives an indication of the I/O rate on the paging file. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Ins Rate The number of kilobytes that the virtual memory manager pages in per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Outs (Per Sec) The average rate per second of page-out requests over a specified sampling period. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Out Rate The number of kilobytes that the virtual memory manager pages out per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Reclaims (Per Sec) The number of times during the monitoring interval that the system removed a page from the queue and used that page for another process. This is the average rate per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Scan Rate (Per Sec) The average rate per second of pages examined over the sampling interval. Virtual memory manager uses a clock-algorithm to implement a pseudo “least recently used (lsu), page replacement scheme. Pages are aged by examination by the clock. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Scanning The number of kilobytes that the virtual memory manager pages scans per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Pages Paged In (Per Sec) The average rate per second of pages that were paged-in from disk to system memory during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Pages Paged Out (Per Sec) The average rate per second of pages paged-out from system memory to disk during the monitoring interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Physical Consumption (AIX) Number of physical CPU units consumed by the LPAR. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Idle Number of processes currently in idle state. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Runnable Number of processes waiting to be run. (For AIX and HP-UX, this includes both the processes able to be run as well as the running processes.) For AIX, they are the processes in the SACTIVE state: active. For HP-UX, they are the processes in the PS_RUN state: running or waiting for CPU. For Solaris they are the processes in the R (SRUN) state: able to be run, but not currently running. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Running Number of processes currently running on a processor (available for Solaris only). For Solaris, it is the number of processes in the O (SONPROC) state: running. This attribute is not available for AIX and HP-UX because, on these systems, this information is collected as part of the Processes Runnable attribute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Sleeping Number of processes currently in sleep state. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Stopped Number of processes currently in the stopped state. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Threads in Run Queue The total number of processes that can be run (or threads in AIX 4 1 and above) waiting for execution by the CPU. This number does not include processes waiting for I/O or some external event, or processes in a sleeping state. The following values are valid: numeric values in the range 0 to 999. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Threads Waiting The number of processes (or threads in AIX 4 1 and above) waiting for page operations. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Processes Zombie Number of zombie processes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Stolen Busy Cycles Pct (AIX) The percentage of physical processor that is comprised of busy cycles stolen by the hypervisor, for dedicated partitions only. Note: the value -100 indicates Not Available, -200 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Stolen Idle Cycles Pct (AIX) The percentage of physical processor that is comprised of idle cycles stolen by the hypervisor, for dedicated partitions only. Note: the value -100 indicates Not Available, -200 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Swap Space Free Amount of swap space (in MB) available. Swap space is usually a disk partition on which page-outs are written. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System Calls Number of system calls made during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System CPU (Percent) The percentage of CPU time devoted to executing UNIX system kernel code. The following values are valid: numeric values expressed as a percentage in the range 1 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

System Name The managed system name.

System Read The number of read and readv system calls during the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System Software Version (AIX) The system software version identification. Valid entry is a simple text string, alphanumeric with a maximum length of 96 characters.

System Write Number of write and writev system calls over the sampling interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Time Spent in Hypervisor Pct (AIX) The percentage of time spent in the hypervisor during the monitoring period. Note: the value -10 indicates Not Available, -20 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Timestamp The date and time the agent collects information as set on the monitored system.

Total Real Memory (KBytes) The total number of kilobytes of physical memory (primary paging memory) on the monitored system. This is a deprecated attribute. New queries should utilize the UNIX Memory attributes for clearer definitions. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Total Virtual Memory (KBytes) The Virtual memory is the total amount memory that is available virtually. The total amount memory that is displayed (virtual) to be available is nothing but the real memory and the part of the disk (paging space) that is available for usage as memory. This is a deprecated attribute. New queries

should utilize the UNIX Memory table attributes for more clear definitions. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Type The type of UNIX operating system residing on a monitored host. For example, you might see AIX, HPUX, SunOS or OSF1 displayed on your system. The following values are valid:

AIX	The IBM AIX operating system
HPUX	The Hewlett Packard HP-UX operating system
SunOS	Sun Microsystems Solaris I or II operating system

- Maximum length 8 characters

Up Time (Seconds) The number of seconds that a monitored system has been running continuously. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

UpTime The system up time of the monitored system. Valid entries in the format DDDdHH:MM:SS, where:

DDD	= Days
HH	= Hours
MM	= Minutes
SS	= Seconds

User CPU (Percent) The percentage of processor time devoted to user processes. The following values are valid: numeric values expressed as a percentage in the range 0 to 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Version The version number of a UNIX operating system on the network. Valid entry is a simple text string, alphanumeric with a maximum length of 16 characters.

Virtual Memory Percent Available Percent of virtual memory available. This is calculated by the agent using the formula:

100 - Percent_Virtual_Memory_Used.

The following values are valid: numeric values in the range 0 to100.0 to one decimal place. Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

Virtual Memory Percent Used Percent of virtual memory used. This calculated by the agent using the formula:

Active_Virtual_Memory / Total_Virtual_Memory * 100.

The following values are valid: numeric values in the range 0 to 100.0 to one decimal place. Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

Wait I/O (Percent) The percentage of time the CPU spends waiting for I/O operations. The following values are valid: numeric values in the range 0 to 100. Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

Zone ID (Solaris) The ID of the Solaris zone. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Zone Name (Solaris) The name of the Solaris zone. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

TCP Statistics Attributes

The TCP Statistics attributes include the number of TCP data packets retransmitted per second, the managed system name, and time stamp. The sampling interval is configurable by means of the KUX_TCP_STAT_SAMPLE_SECS environment variable. The default value is 30 seconds, and the minimum value is 5 seconds.

System Name The managed system name.

TCP Data Packets Retransmitted Per Second The number of TCP data packets retransmitted per second. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Time Stamp The date and time the agent collects information as set on the monitored system.

UNIX Memory attributes

The UNIX Memory attributes refer to memory characteristics.

The use of the term megabyte refers to 2^20 bytes. This is equivalent to the SI unit, megabyte.

Available Real Memory (MB) The amount of physical memory, in megabytes, currently available on the system. The amount of real memory available to user programs and the system. It is not unusual for the amount of available memory to be low, because the system uses available real memory for caching of system information. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Note: This is the amount of physical RAM available for program allocation. This value is derived from the number of free pages available in real memory.

Available Swap Space (MB) The number of megabytes of secondary storage available to host virtual memory. Any part of memory available for additional paging (pseudo-swap) that is free is also included. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Note: This value represents the amount of paging space the system has available to assign programs memory.

Available Virtual Storage (MB) The total number of megabytes available for paging that is not currently being used. This includes primary memory and secondary storage. This includes primary memory and secondary storage. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Note: The amount of memory available to the system for allocation. It is not unusual for this value to be small on some operating systems because the system uses real memory as a caching space.

Average Page Ins\1 Minute The average rate of page ins over the last 1 minute. Displays N/C (not collected) if agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Page Ins\5 Minutes The average rate of page ins over the last 5 minutes. Displays N/C (not collected) if agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Page Ins\15 Minutes The average rate of page ins over the last 15 minutes. Displays N/C (not collected) if agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Page Ins\60 Minutes The average rate of page ins over the last 60 minutes. Displays N/C (not collected) if agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Page Outs\1 Minute The average rate of page outs over the last 1 minute. Displays N/C (not collected) if agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Page Outs\5 Minutes The average rate of page outs over the last 5 minutes. Displays N/C (not collected) if agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Page Outs\15 Minutes The average rate of page outs over the last 15 minutes. Displays N/C (not collected) if agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Page Outs\60 Minutes The average rate of page outs over the last 60 minutes. Displays N/C (not collected) if agent has been up for less than 60

minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Page Scans\1 Minute The average rate of page scans over the last 1 minute. Displays N/C (not collected) if agent has been up for less than 1 minute. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Page Scans\5 Minutes The average rate of page scans over the last 5 minutes. Displays N/C (not collected) if agent has been up for less than 5 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Page Scans\15 Minutes The average rate of page scans over the last 15 minutes. Displays N/C (not collected) if agent has been up for less than 15 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Average Page Scans\60 Minutes The average rate of page scans over the last 60 minutes. Displays N/C (not collected) if agent has been up for less than 60 minutes. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Computational Memory (AIX) The number of computational 4K pages resident in memory. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Decay Rate (AIX) The decay rate for repaging values per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Net Memory Available (MB) The number of megabytes of physical memory available on the system considering as free memory the amount of memory used by ZFS Adaptive Replacement Cache. This information is available just on Solaris machines. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum.

Net Memory Available (Percent) The percent of physical memory available on the system considering as free memory the amount of memory used by ZFS Adaptive Replacement Cache. This information is available just on Solaris machines. Note: the value -1 indicates Not Available, -2 indicates Not Collected.

Net Memory Used (MB) The number of megabytes of physical memory free on the system considering as free memory the amount of memory used by ZFS Adaptive Replacement Cache. This information is available just on Solaris machines. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum.

Net Memory Used (Percent) The percent of physical memory used on the system considering as free memory the amount of memory used by ZFS Adaptive

Replacement Cache. This information is available just on Solaris machines. Note: the value -1 indicates Not Available, -2 indicates Not Collected.

Non Computational Memory (AIX) The number of non-computational 4K pages resident in memory. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Faults Per Second Page faults per second averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page In Rate (KB per Second) The number of kilobytes that the virtual memory manager pages in per second averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page-In Requests Per Second The number of requests per second of page-in requests averaged over the previous 30-second interval. A page-in request may include multiple pages and gives an indication of the I/O rate on the paging file. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Out Rate (KB per Second) The number of kilobytes that the virtual memory manager pages out per second averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page-Out Requests Per Second The number of requests per second of page-out requests averaged over the previous 30-second interval. A page-out request may include multiple pages and gives an indication of the I/O rate on the paging file. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Reclaims Per Second The number of times per second that the system removed a page from the queue and used that page for another process, averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Scan Rate (KB Per Second) The number of kilobytes in pages that the virtual memory manager scans per second averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Page Scan Rate Per Second The number of pages examined per second over the previous 30-second interval. Virtual memory manager uses a clock-algorithm to implement a pseudo least recently used (lsu), page replacement scheme. Pages are aged by examination by the clock. Note: the value -1 indicates Not Available, -2

indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Pages Paged-In Per Second The number of pages per second that were paged-in from disk to system memory averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Pages Paged-Out Per Second The number of pages per second paged-out from system memory to disk averaged over the previous 30-second interval. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Pages Read per Sec (AIX) The number of 4K pages read by VMM per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Pages Written per Sec (AIX) The number of 4K pages written by VMM per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Paging Space Free Pct (AIX) The percentage of system paging space that is free. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Paging Space Read per Sec (AIX) The number of 4K pages read from paging space by VMM per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Paging Space Used Pct (AIX) The percentage of system paging space that is used. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Paging Space Written per Sec (AIX) The number of 4K pages written to paging space by VMM per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Percent Real Memory Available Percent of real memory available. Calculated by the agent using the formula: (Avail_Real_Mem_MB / Total_Real_Mem_MB) * 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Percent Real Memory Used Percent of real memory used. Calculated by the agent using the formula: (100 - Avail_Real_Mem_Pct). Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

Percent Swap Space Available Percent of swap space available. Calculated by the agent using the formula: (Avail_Swap_Space_MB / Total_Swap_Space_MB) * 100. Note: the value -1 indicates Not Available, and -2 indicates Not Collected.

Percent Swap Space Used The percent of swap space used. Calculated by the agent using the formula: (100 - Avail_Swap_Space_Pct). Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Percent Virtual Storage Available Percent of virtual storage available. Calculated by the agent using the formula: 100 - Virtual_Storage_Pct_Used. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Percent Virtual Storage Used Percent of virtual memory used. Calculated by the agent using the formula: (Used_Virtual_Storage_MB / Total_Virtual_Storage_MB) * 100. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Repaging Rate (AIX) The global repaging rate per second. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

System Name The managed system name.

Timestamp The date and time the agent collects information as set on the monitored system.

Total Real Memory (MB) The total number of megabytes of physical memory on a monitored system. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Note: The total physical RAM installed and active in the system. Some systems have the capability to disable portions of RAM; the disabled memory is not reported. On virtual machines, the memory reported is that which is allocated to the virtual machine and activated by the operating system.

Total Swap Space (MB) The total number of megabytes of secondary storage dedicated to hosting virtual memory. Any part of memory available for the system to use for additional paging (pseudo-swap) is also included. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Note: The total amount of executable pages possible for paging including physical RAM plus secondary paging space, in megabytes. All of the physical RAM might not be available for paging, therefore, total swap space might not equal total virtual storage. Total swap space does not include file cache on AIX systems.

Total Virtual Storage (MB) The total number of megabytes of storage available for hosting virtual memory. This includes primary memory and secondary backing store. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Note: Total virtual storage is the total physical RAM allocated to the system plus the total secondary paging space. Total virtual storage combines the total real RAM plus the total secondary storage.

Used Real Memory (MB) The amount of physical memory, in megabytes, currently used on the system. The amount of real memory in use by user programs and the system. It is not unusual for the amount of in use memory to be very high,

because the system uses available real memory for caching of system information. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Note: Used Real Memory is derived from subtracting Available Real Memory from Total Real Memory. Used Real Memory includes the following:

- Executable Pages stored in real memory
- System Pages stored in real memory
- Application Heap Pages stored in real memory
- Shared Memory Pages stored in real memory
- MMAP pages stored in real memory
- File Cache pages
- Memory that is outside of Page Space, usually used by the kernel

Used Swap Space (MB) The number of megabytes of secondary storage currently hosting virtual memory. Any part of memory used for additional paging (pseudo-swap) is also included. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Note: The total amount of pages being used by the system in real memory and secondary paging space. This memory includes any physical memory being used as pseudo-swap and reserved memory. Reserved memory includes allocations that have been allocated, but not accessed and therefore not assigned to physical memory or secondary paging spaces. Used swap space consists of the following:

- Executable Pages stored in real memory and on secondary paging space
- System Pages stored in real memory and secondary paging space
- Application Heap Pages stored in real memory and secondary paging space
- Shared Memory Pages stored in real memory and secondary paging space
- MMAP pages stored in real memory and secondary paging space

Used Virtual Storage (MB) The number of megabytes of virtual memory currently in use by the system. This includes primary memory and secondary backing store. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Note: All the secondary paging space plus all the real RAM used for any purpose. This could include all memory included in used real memory and that allocated on the disk.

ZFS ARC Size (MB) The number of megabytes of physical memory used on the system by ZFS Adaptive Replacement Cache. This information is available just on Solaris machines. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum.

User attributes

The User Attributes refer to user characteristics such as idle time, user name, location, and login time.

Idle Time The number of minutes that have passed since a user last entered a command. Use this attribute to check idle time. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Location Information provided by the user about their location. Valid entry is a simple text string, alphanumeric with a maximum length 16 characters. This information varies from location to location, and might not be available for all users or for all UNIX operating systems. Check with your local System Administrator for additional information concerning this attribute.

Login Name The login name of a user. Valid entry is a simple text string, alphanumeric with a maximum length 16 characters. Use this attribute to include or exclude specific user login names in the situation.

Login Name (Unicode) The login name of a user. Use this attribute to include or exclude specific user login names in the situation.

Login Time The date and time a user logged in. Valid entries are in the format CYYMMDDHHMMSSmmm, where:

C	= Century (0=20th)
YY	= Year
MM	= Month of the Year (01-12)
DD	= Day of the Month (01-31)
HH	= Hour, in 24-hour time (00-23)
MM	= Minute
SS	= Second
mmm	= Millisecond

Example To express November 6, 1998, 1:05 p.m., enter 0981106130500000.

Name The full name of a user. Valid entry is a simple text string, alphanumeric with a maximum length 32 characters.

Name (Unicode) The full name of a user.

System Name The managed system name.

Terminal The identity of a logged-in device. Valid entry is a simple text string, alphanumeric with a maximum length 8 characters.

Example A value of w14921 could be the name of a device. This could indicate a location, such as Wood Lake. Your location might already have a naming convention for network devices. Check with your local system administrator for additional information.

Timestamp The date and time the agent collects information as set on the monitored system.

User ID The numeric ID the system assigned to a user. Note: the value -1 indicates Not Available, -2 indicates Not Collected, 2147483647 indicates Value_Exceeds_Maximum, and -2147483648 indicates Value_Exceeds_Minimum.

Use this attribute to include or exclude a particular user in the situation. This attribute is supported on all platforms. On AIX, you need a patch from IBM in order to get the user ID.

The numeric identification number varies from system to system and user to user. An example of a user ID could be 48765.

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 with historical data that is 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.

Expected number of instances 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 on the application environment that is being monitored. For example, if your attribute group is monitoring each processor on your computer and you have a dual processor computer, the number of instances is 2.

Calculate expected disk space consumption by multiplying the number of bytes per instance by the expected number of instances, and then multiplying that product by the number of samples. Table 2 provides the following information required to calculate disk space for the Monitoring Agent for UNIX OS:

- *Bytes per instance (agent)* 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.
- *Database bytes per instance (warehouse)* is an estimate of the record length for detailed records written to the warehouse database, if the attribute group is configured to be written to the warehouse. Detailed records are those that have been uploaded from the agent for long-term historical data collection. This estimate can be used for warehouse disk space planning purposes.
- *Aggregate bytes per instance (warehouse)* is an estimate of the record length for aggregate records written to the warehouse database, if the attribute group is configured to be written to the warehouse. Aggregate records are created by the Summarization agent for attribute groups that have been configured for summarization. This estimate can be used for warehouse disk space planning purposes.
- *Number of rows stored in warehouse per historical collection query* is an estimate of the number of rows that are stored in the Tivoli Data Warehouse per query.

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 2. Capacity planning for historical data logged by component

Table	Attribute group	Bytes per instance (agent)	Database bytes per instance (warehouse)	Aggregate bytes per instance (warehouse)
UNIXDISK	`Disk`	1212	1285	2166
UNIXDPERF	`Disk_Performance`	328	299	843
FILEINFO	`File_Information`	4212	4260	4508

Table 2. Capacity planning for historical data logged by component (continued)

Table	Attribute group	Bytes per instance (agent)	Database bytes per instance (warehouse)	Aggregate bytes per instance (warehouse)
KUXPASSTAT	KCA_UX_Agent_Active_Runtime_Status	1410	1435	1472
KUXPASMGMT	KCA_UX_Agent_Availability_Management_\ Status	538	542	579
KUXPASALRT	KCA_UX_Alerts_Table	512	516	553
KUXPASCAP	KCA_UX_Configuration_Information	3090	3132	3169
UNIXMACHIN	Machine_Information	536	540	577
UNIXNFS	N_F_S_and_R_P_C_Statistics	520	449	3138
UNIXNET	Network	524	605	2022
UNIXPS	Process	2288	2264	3003
UNIXCPU	SMP_CPU	364	264	1471
UNIXSOLZON	Solaris_Zones	550	583	800
UNIXOS	System	844	846	3670
UNIXIPADDR	UNIX_IP_Address	574	578	615
UNIXALLUSR	Unix_All_Users	188	189	265
UNIXFILCMP	Unix_File_Comparison	1652	1660	1697
UNIXFILPAT	Unix_File_Pattern	1652	1660	1697
UNIXGROUP	Unix_Group	164	163	200
UNIXMEM	Unix_Memory	256	361	1874
UNIXPING	Unix_Ping	884	899	987
UNXPRINTQ	Unix_Print_Queue	316	309	424
UNIXUSER	User	312	310	386

Note: FILEINFO cannot be historically collected.

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

Chapter 5. Situations reference

This chapter contains an overview of situations, references for detailed information about situations, and descriptions of the predefined situations included in this monitoring agent.

About situations

A situation is a logical expression involving one or more system conditions. Situations are used to monitor the condition of systems in your network. You can manage situations from the Tivoli Enterprise Portal by using the Situation editor.

The IBM Tivoli Monitoring agents that you use to monitor your system environment are delivered with a set of predefined situations that you can use as-is or you can create new situations to meet your requirements. Predefined situations contain attributes that check for system conditions common to many enterprises.

Using predefined situations can improve the speed with which you can begin using the Monitoring Agent for UNIX OS. You can examine and, if necessary, change the conditions or values being monitored by a predefined situation to those best suited to your enterprise.

Note: The predefined situations provided with this monitoring agent are not read-only. Do not edit these situations and save over them. Software updates will write over any of the changes that you make to these situations. Instead, clone the situations that you want to change to suit your enterprise.

You can display predefined situations and create your own situations using the Situation editor. The left frame of the Situation editor initially lists the situations associated with the Navigator item that you selected. When you click a situation name or create a new situation, the right frame opens with the following tabs:

Formula
Condition being tested

Distribution
List of managed systems (operating systems, subsystems, or applications) to which the situation can be distributed.

Expert Advice
Comments and instructions to be read in the event workspace

Action
Command to be sent to the system

Until Duration of the situation

The Monitoring Agent for UNIX OS predefined situations describe system conditions on your UNIX networked systems that you want to monitor at your site. Use these situations to being monitoring any UNIX managed object quickly, or as models for customizing your own situations. In some cases, the values that are assigned to the predefined situations are examples only and should be modified to reflect the conditions of your distributed system. Each predefined situation is assigned to a predefined template, and an alert status for the situation is defined.

More information about situations

The *IBM Tivoli Monitoring User's Guide* contains more information about predefined and custom situations and how to use them to respond to alerts.

For a list of the predefined situations for this monitoring agent and a description of each situation, refer to the Predefined situations section below and the information in that section for each individual situation.

Predefined situations

This monitoring agent contains the following predefined situations, which are organized by the Navigator node that the situations are associated with.

- Agent Management Services node
  - UMX_AMS_Alert_Critical
- Disk Usage node
  - UNIX_AIX_Avg_ReqInWaitQ_MS_Info
  - UNIX_AIX_Avg_Transfer_MS_Info
  - UNIX_AIX_ServQ_Full_PerSec_Info
  - UNIX_BP_SpaceUsedPct_Critical
  - UNIX_BP_SpaceUsedPctCustom_Crit
  - UNIX_CMD_Disk_Inodes_Critical
  - UNIX_CMD_Disk_Inodes_Critical_2
  - UNIX_CMD_Disk_Space_Warning
  - UNIX_Disk_Availability
  - UNIX_Filemount_Critical
  - UNIX_HD_Config_Critical
  - UNIX_HD_Config_Critical_2
  - UNIX_scratch_tmp_Disk_Full
- File Information node
  - UNIX_User_File_Exists
- Network node
  - UNIX_AIX_NetBandwidth_High_Info
  - UNIX_BP_NetInOutErrPct_Critical
  - UNIX_Network_Collsns_Critical
  - UNIX_Network_Collsns_Warning
  - UNIX_Network_Errors
  - UNIX_Network_Interface_Busy
  - UNIX_Network_Interface_Idle
- NFS Activity node
  - UNIX_NFS_RPC_Rejects
- Process node
  - UNIX_AIX_Process_ResDat_Hi_Info
  - UNIX_AIX_Process_ResTxt_Hi_Info
  - UNIX_BP_ProcHighCpu_Critical
  - UNIX_BP_ProcMissing_Critical
  - UNIX_CMD_Process_Critical

  - UNIX_CMD_Runaway_Process
  - UNIX_CPU_Critical
  - UNIX_CPU_Warning
  - UNIX_Process_Memory_Critical
  - UNIX_Process_Memory_Leak
  - UNIX_Process_Memory_Warning
  - UNIX_Process_MISSING_inetd
- System Information node
  - UNIX_Active_Virtual_Memory
  - UNIX_AIX_CPU_CtxSwitch_Hi_Info
  - UNIX_AIX_Device_Stopped_Warning
  - UNIX_AIX_Memory_RePg_Hi_Info
  - UNIX_AIX_System_HypPct_Hi_Info
  - UNIX_AIX_System_NProcs_Hi_Info
  - UNIX_AIX_User_Acct_Locked_Info
  - UNIX_AIX_User_Login_Retry_Info
  - UNIX_BP_AvgCpuBusyPct5min_Criti
  - UNIX_BP_CpuBusyPct_Critical
  - UNIX_BP_LoadAvg5min_Critical
  - UNIX_BP_NumberZombies_Warning
  - UNIX_BP_PagingRate_Critical
  - UNIX_BP_SwapSpaceUsedPct_Critic
  - UNIX_BP_SysWaitIOPct_Warning
  - UNIX_CPU_Busy_Critical
  - UNIX_CPU_Busy_Warning
  - UNIX_HD_Excessive_IO_Wait
  - UNIX_LPARBusy_pct_Warning
  - UNIX_LPARPhyBusy_pct_Warning
  - UNIX_LPARvcs_Info
  - UNIX_LPARfreepool_Warning
  - UNIX_LPARPhanIntrs_Info
  - UNIX_LPARentused_Info
  - UNIX_LPAR_MaxCPUCapUsed_Info
  - UNIX_LPAR_Moved_Info
  - UNIX_System_Busy_Critical
  - UNIX_System_Busy_Warning
  - UNIX_System_Capacity_Critical
  - UNIX_System_Paging_Critical
  - UNIX_System_Virtual_Memory_Warning
  - UNIX_User_CPU_Critical
  - UNIX_WPAR_Admin_Op_Info
  - UNIX_WPAR_Broken_Warning
  - UNIX_WPAR_CPU_Usage_Warning
  - UNIX_WPAR_Mem_Usage_Warning
  - UNIX_WPAR_Min_CPU_Limit_Info

- UNIX_WPAR_Min_Mem_Limit_Info
- UNIX_WPAR_RC_Inactive_Info
- UNIX_WPAR_Unlim_CPU_Shares_Info
- UNIX_WPAR_Unlim_Mem_Shares_Info

The individual predefined situations below are listed under the workspace associated with the situation.

Agent Management Services node

UNIX_AMS_Alert_Critical situation

Determines if one of the following conditions is true:

- A managed agent has exceeded its restart count for the day as configured in the 'maxRestarts' field of its Common Agent Package file.
- A managed agent is overutilizing the available CPU resources as configured in the 'cpuThreshold' field of its Common Agent Package file.
- A managed agent is overutilizing the available system memory resources as configured in the 'memoryThreshold' field of its Common Agent Package file.
- An attempt at auto-restarting a managed agent failed.
- An attempt at starting a stopped or manually stopped managed agent failed.
- The Agent Management Services watchdog is no longer reliable. If either watchdog stops monitoring, you will receive this message.

The formula for this situation is as follows:

```
Alert Message=='Agent exceeded restart count' OR
Alert Message=='Agent overutilizing CPU' OR
Alert Message=='Agent overutilizing memory'  OR
Alert Message=='Agent restart failed' OR
Alert Message=='Agent manual stop failed' OR
Alert Message =='Agent Management Services watchdog no longer reliable'
```

Disk Usage node

UNIX_AIX_Avg_ReqInWaitQ_MS_Info situation

Monitors if the average time waiting for disk access is high.

The formula for this situation is as follows:

```
Disk_Performance.Avg_Wait GT 20
```

UNIX_AIX_Avg_Transfer_MS_Info situation

Monitors if the average amount of disk time used is high.

The formula for this situation is as follows:

```
Disk_Performance.Avg_Serv GT 5
```

UNIX_AIX_ServQ_Full_PerSec_Info situation

Monitors if the number of times the service queue becomes full per second is high.

The formula for this situation is as follows:

```
Disk_Performance.ServiceQ_Full_per_Sec GT 5
```

UNIX_BP_SpaceUsedPct_Critical situation

Monitors all mounted file systems for space used percentage. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

The formula for this situation is as follows:

```
IF VALUE
Disk.Space_Used_Percent GE 95
```

UNIX_BP_SpaceUsedPctCustom_Crit situation

Monitors only specific file systems for space used percentage, for example /home. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

The formula for this situation is as follows:

```
IF ( ( VALUE Disk.Mount_Point_U EQ '/'
AND VALUE Disk.Space_Available_Percent LT 10 )
OR ( VALUE Disk.Mount_Point_U EQ '/home'
AND VALUE Disk.Space_Available_Percent LT 20 ) )
UNTIL ( SIT UNIX_BP_SpaceUsedPct_Critical )
```

UNIX_CMD_Disk_Inodes_Critical situation

Superseded by UNIX_CMD_Disk_Inodes_Critical_2. Monitors the /tmp and /var free inodes for critical space.

The formula for this situation is as follows:

```
IF Disk.Mount_Point EQ /tmp OR Disk.Mount_Point  EQ /var) AND Disk.Inodes_Free
LT 20000 ACTION echo UNIX_CMD_Disk_Inodes_Critical  &Disk.System_Name Low free
inodes on /tmp and /var
```

UNIX_CMD_Disk_Inodes_Critical_2 situation

Monitors the /tmp and /var free inodes for critical space.

The formula for this situation is as follows:

```
IF Disk.Mount_Point_U EQ /tmp OR Disk.Mount_Point_U  EQ /var)
AND Disk.Inodes_Free_64
LT 20000 ACTION echo UNIX_CMD_Disk_Inodes_Critical_2
&Disk.System_Name Low free inodes on /tmp and /var
```

UNIX_CMD_Disk_Space_Warning situation

Monitors any mounted file system with space usage greater than 90 percent.

The formula for this situation is as follows:

```
IF Disk.Space_Used_Percent GE 90 ACTION  echo UNIX_CMD_Disk_Space_Warning
&Disk.System_Name Filemount: &Disk.Mount_Point  Space_Used: &Disk.Space_Used_Percent
```

UNIX_Disk_Availability situation

Determines under-utilized hard disk space.

The formula for this situation is as follows:

```
IF Disk.Mount_Point EQ /user AND Disk.Space_Used_Percent  LT 25
```

UNIX_Filemount_Critical situation

Checks for the existence of a specific mount point on a specific system.

The formula for this situation is as follows:

```
IF Disk.System_Name EQ Redwood AND Disk.Mount_Point  EQ /usr
```

UNIX_HD_Config_Critical situation

Superseded by UNIX_HD_Config_Critical_2. Monitors hard disk space or free inodes that are going critical.

The formula for this situation is as follows:

```
IF Disk.Inodes_Free LT 100 OR Disk.Space_Used_Percent  GT 90
```

UNIX_HD_Config_Critical_2 situation

Monitors hard disk space or free inodes that are going critical.

The formula for this situation is as follows:

```
IF Disk.Inodes_Free_64 LT 100 OR Disk.Space_Used_Percent  GT 90
```

UNIX_scratch_tmp_Disk_Full situation

Monitors file mount /scratch or /tmp with space usage greater than 90 percent.

The formula for this situation is as follows:

```
IF Disk.Space_Used_Percent GT 90 AND (SCAN  Disk.Mount_Point EQ /scratch OR
Disk.Mount_Point EQ /tmp)
```

File Information node

UNIX_User_File_Exists situation

Monitors for the existence of a specific user file.

The formula for this situation is as follows:

```
IF File_Information.Path EQ /a/path2/search  AND File_Information.File EQ
the File_2find
```

Network node

UNIX_AIX_NetBandwidth_High_Info situation

Monitors if the percentage of physical network adapter bandwidth utilization is high.

The formula for this situation is as follows:

```
Network.Bandwidth_Util_Pct GT 60
```

UNIX_BP_NetInOutErrPct_Critical situation

Monitors the percentage of errors on received or transmitted network packets. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

The formula for this situation is as follows:

```
IF (( VALUE Network.Interface_Status EQ Up AND VALUE Network.Input_
Packet_Errors_PercentGT 10 ) OR ( VALUE Network.Interface_Status
EQ Up AND VALUENetwork.Output_Packet_Errors_Percent GT 10 ) )
```

UNIX_Network_Collsns_Critical situation

Indicates a large number of network collisions.

The formula for this situation is as follows:

```
IF Network.Collisions GT 15
```

UNIX_Network_Collsns_Warning situation
Indicates a small number of network collisions.

The formula for this situation is as follows:

```
IF Network.Collisions GT 2
```

UNIX_Network_Errors situation
Monitors whether the received or transmitted error limit has been exceeded.

The formula for this situation is as follows:

```
IF Network.Interface_Status EQ UP AND (Network.Output_Errors  GT 10
OR Network.Input_Errors GT 10)
```

UNIX_Network_Interface_Busy situation
Monitors whether the frames transmitted or received has exceeded the limit.

The formula for this situation is as follows:

```
IF Network.Network_Interface_Name NE Lo0  AND Network.Interface_Status EQ
UP AND (Network.Frames_Received GT 1000  OR Network.Frames_Transmitted GT
1000)
```

UNIX_Network_Interface_Idle situation
Monitors whether the frames transmitted or received is less than the limit.

The formula for this situation is as follows:

```
IF Network.Network_Interface_Name NE Lo0  AND Network.Interface_Status EQ
UP AND (Network.Frames_Received LT 100  OR Network.Frames_Transmitted LT 100)
```

NFS Activity workspace

UNIX_NFS_RPC_Rejects situation
Monitors for rejected NFS RPC calls.

The formula for this situation is as follows:

```
IF N_F_S_and_R_P_C_Statistics.NFS_Server_Calls_Rejected  GT 2
OR N_F_S_and_R_P_C_Statistics.NFS_Client_Calls_Rejected GT 2
```

Process node

UNIX_AIX_Process_ResDat_Hi_Info situation
Monitors if the amount of resident physical memory used by the process private data (4K pages) is high.

The formula for this situation is as follows:

```
Process.Resident_Data_Size GT 100
```

UNIX_AIX_Process_ResTxt_Hi_Info situation
Monitors if the amount of resident physical memory used by the process code (4K pages) is high.

The formula for this situation is as follows:

```
Process.Resident_Text_Size GT 100
```

UNIX_BP_ProcHighCpu_Critical situation

Monitors the CPU percent utilization by all processes except kproc, swapper and wait. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

The formula for this situation is as follows:

```
IF VALUE Process.CPU_Pct GE 95
AND SCAN Process.Process_Command_U NE 'wait'
AND SCAN Process.Process_Command_U NE 'swapper'
AND SCAN Process.Process_Command_U NE 'kproc'
```

UNIX_BP_ProcMissing_Critical situation

Monitors specified processes that are not found in the system. The process might have been killed or might never have been started. An asterisk is used to identify the specific process started from the system directory. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

The formula for this situation is as follows:

```
IF MISSING Process.Command_U EQ ( '/usr//cron','/usr/*/dsmcad','
/usr/*/gmond','/usr/*/inetd','/usr/*/ntpd','/usr/*/snmpd','/usr/*/sshd',
'/usr/*/syslogd' )
```

UNIX_CMD_Process_Critical situation

Monitors for the existence of a process.

The formula for this situation is as follows:

```
IF Process.Command EQ FOO
```

UNIX_CMD_Runaway_Process situation

Reports processes with high CPU utilization.

The formula for this situation is as follows:

```
IF Process.CPU_Utilization GT 95 AND Process.User_ID  NE 0 AND
(Process.Execution_State EQ R OR Process.Execution_State EQ A)
ACTION echo UNIX_CMD_Runaway_Process &Process.System_Name
Processid:  &Process.Process_ID Command: &Process.Command
```

UNIX_CPU_Critical situation

Monitors for processes with CPU utilization that is greater than or equal to 85 times a process uses the CPU over a period of 20 system clock ticks.

The formula for this situation is as follows:

```
IF Process.CPU_Utilization GE 85 AND Process.Command NE kproc AND Process.Command
NE swapper
```

UNIX_CPU_Warning situation

Monitors for processes with CPU utilization that is greater than or equal to 70 and less than 85 times a process uses the CPU over a period of 20 system clock ticks.

The formula for this situation is as follows:

```
IF Process.CPU_Utilization GE 70 AND Process.CPU_Utilization LT 85
```

UNIX_Process_Memory_Critical situation

Reports process with high memory usage that have reached a critical state.

The formula for this situation is as follows:

```
IF Process.Mem_Pct GT 8000
```

UNIX_Process_Memory_Leak situation

Reports process with high virtual memory usage.

The formula for this situation is as follows:

```
IF Process.Virtual_Size GT 9999999
```

UNIX_Process_Memory_Warning situation

Reports processes with high memory usage before they become critical.

The formula for this situation is as follows:

```
IF Process.Mem_Pct GT 5000 AND Process.Mem_Pct  LT 8000
```

UNIX_Process_MISSING_inetd situation

Monitors whether the inetd Internet services daemon is up and running.

The formula for this situation is as follows:

```
IF MISSING Process.Command EQ ('inetd')
```

System Information node

UNIX_Active_Virtual_Memory situation

Monitors whether active virtual memory is approaching total virtual memory.

The formula for this situation is as follows:

```
IF System.Active_Virtual_Memory GE nnnn
```

UNIX_AIX_CPU_CtxSwitch_Hi_Info situation

Monitors if the number of CPU context switches per second is high.

The formula for this situation is as follows:

```
SMP_CPU.Context_Switches_per_Sec GT 1000
```

UNIX_AIX_Device_Stopped_Warning situation

Monitors if the state of an AIX device is stopped.

The formula for this situation is as follows:

```
AIX_DEVICES.State EQ 'Stopped'
```

UNIX_AIX_Memory_RePg_Hi_Info situation

Monitors if the global repaging rate per second is high.

The formula for this situation is as follows:

```
Unix_Memory.Repaging_Rate GT 10
```

UNIX_AIX_System_HypPct_Hi_Info situation

Monitors if the time spent in the hypervisor during the monitoring period is greater than 3%.

The formula for this situation is as follows:

```
System.Time_Spent_in_Hypervisor_Pct GT 3
```

UNIX_AIX_System_NProcs_Hi_Info situation
Monitors if the total number of processes is high.

The formula for this situation is as follows:

```
System.Total_Num_Processes GT 1000
```

UNIX_AIX_User_Acct_Locked_Info situation
Monitors if the user account has been locked.

The formula for this situation is as follows:

```
AIX_Defined_Users.Account_Locked EQ 'true'
```

UNIX_AIX_User_Login_Retry_Info situation
Monitors if the total number of login retries is high.

The formula for this situation is as follows:

```
AIX_Defined_Users.Loginretries GT 4
```

UNIX_BP_AvgCpuBusyPct5min_Criti situation
Monitors the total percentage of CPU (system + user) busy for the sum of all the CPU on the system. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

Note: This situation is not applicable to single CPU systems.

The formula for this situation is as follows:

```
IF VALUE SMP_CPU.Avg_CPU_Busy_5 GT
90 AND VALUE SMP_CPU.CPU_ID  EQ aggregate
```

UNIX_BP_CpuBusyPct_Critical situation
Monitors if the CPU workload is high. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

The formula for this situation is as follows:

```
IF ( ( VALUE SMP_CPU.CPU_Busy
GT 90 AND VALUE SMP_CPU.CPU_ID E Q aggregate ) OR ( VALUE SMP_CPU.CPU_Busy
GT 95 AND VALUE SMP_CPU.C PU_ID NE aggregate ) )
```

UNIX_BP_LoadAvg5min_Critical situation
Monitors the average number of processes in the UNIX kernel run queue during the last five minutes. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

The formula for this situation is as follows:

```
IF ( ( VALUE Machine_Information.Number_of_Physical_Processors EQ 1
AND VALUE System.Load_Average_5_Min GT 4.0 )
OR ( VALUE Machine_Information.Number_of_Physical_Processors EQ 2
AND VALUE System.Load_Average_5_Min GT 8 )
OR ( VALUE Machine_Information.Number_of_Physical_Processors EQ 3
AND VALUE System.Load_Average_5_Min GT 12 )
OR ( VALUE Machine_Information.Number_of_Physical_Processors EQ 4
AND VALUE System.Load_Average_5_Min GT 16 )
OR ( VALUE Machine_Information.Number_of_Physical_Processors EQ 6
AND VALUE System.Load_Average_5_Min GT 24 )
OR ( VALUE Machine_Information.Number_of_Physical_Processors EQ 8
```

```
AND VALUE System.Load_Average_5_Min GT 32 )
OR ( VALUE Machine_Information.Number_of_Physical_Processors
GE 16 AND VALUE System.Load_Average_5_Min GT 64 ))
```

UNIX_BP_NumberZombies_Warning situation

Monitors the number of defunct processes. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

The formula for this situation is as follows:

```
IF VALUE System.Processes_Zombie
GE 50
```

UNIX_BP_PagingRate_Critical situation

Monitors the average rate per second of pages examined over the sampling interval. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

The formula for this situation is as follows:

```
IF VALUE System.Page_Scan_Rate GT 500
```

UNIX_BP_SwapSpaceUsedPct_Critic situation

Monitors the percentage of swap space used. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

The formula for this situation is as follows:

```
IF VALUE Unix_Memory.Used_Swap_Space_Pct
GT 40.0
```

UNIX_BP_SysWaitIOPct_Warning situation

Monitors the percent of time the CPU spends waiting for I/O operations. This situation is automatically installed, but not automatically distributed, during installation to the default MSL or managed system.

The formula for this situation is as follows:

```
IF
VALUE System.Wait_I/O GT 20
```

UNIX_CPU_Busy_Critical situation

Monitors whether the CPU workload is high (greater than 90 percent).

The formula for this situation is as follows:

```
IF SMP_CPU.CPU_Busy GT 90
```

UNIX_CPU_Busy_Warning situation

Monitors whether the CPU workload is greater than 70 percent and less than or equal to 90 percent.

The formula for this situation is as follows:

```
IF SMP_CPU.CPU_Busy GT 70 AND SMP_CPU.CPU_Busy  LE 90
```

UNIX_HD_Excessive_IO_Wait situation

Monitors a typical I/O bound processor (NSF).

The formula for this situation is as follows:

```
IF System.Wait_I/O GT 20
```

UNIX_LPARBusy_pct_Warning situation

Monitors if the logical busy time is greater than 95%.

The formula for this situation is as follows:

```
AIX_LPAR.Busy_Pct GT 95
```

UNIX_LPARPhyBusy_pct_Warning situation

Monitors if the physical busy time of a full processor is greater than 95%.

The formula for this situation is as follows:

```
AIX_LPAR.Phys_Busy_Pct GT 95
```

UNIX_LPARvcs_Info situation

Monitors if the virtual CPU context switches is greater than 1000 per second.

The formula for this situation is as follows:

```
AIX_LPAR.Virt_Context_CPU_Switches_per_Sec GT 1000
```

UNIX_LPARfreepool_Warning situation

Monitors if the unallocated capacity available in the shared pool is less than 100.

The formula for this situation is as follows:

```
AIX_LPAR.Unallocated_CPU_In_Pool LT 100
```

UNIX_LPARPhanIntrs_Info situation

Monitors if the number of phantom interrupts is greater than 100.

The formula for this situation is as follows:

```
AIX_LPAR.Phantom_Interrupts GT 100
```

UNIX_LPARentused_Info situation

Monitors if entitlement is greater than 100%.

The formula for this situation is as follows:

```
AIX_LPAR.Entitlement_Pct GT 100
```

UNIX_LPAR_MaxCPUCapUsed_Info situation

Monitors if the percentage of maximum physical CPU available to this LPAR that was actually used is greater than 80%.

The formula for this situation is as follows:

```
AIX_LPAR.Max_CPU_Cap_Used_Pct GT 80
```

UNIX_LPAR_Moved_Info situation

Monitors if the frame hardware ID to which this LPAR belongs is not equal to the previous frame hardware ID of this LPAR before it was migrated to the current frame.

The formula for this situation is as follows:

```
AIX_LPAR.Machine_ID NE AIX_LPAR.Last_Machine_ID
```

UNIX_System_Busy_Critical situation
Monitors for a critical state of I/O wait, low free memory, and CPU idle.

The formula for this situation is as follows:

```
IF System.Wait_I/O GT 25 AND System.Free_Memory  LT 1 AND System.Idle_CPU
GT 10 AND System.Load_Average_1_Min GT 2
```

UNIX_System_Busy_Warning situation
Monitors for system CPU, idle, I/O wait, and load average for busy state.

The formula for this situation is as follows:

```
IF System.System_CPU GT 50 AND System.Idle_CPU  GT 0 AND System.Wait_I/O GT
0 AND System.Load_Average_5_Min GT 1
```

UNIX_System_Capacity_Critical situation
Monitors system capacity using a process number and CPU usage.

The formula for this situation is as follows:

```
IF System_Proc_Number GE 250 AND System.System_CPU  GT 80
```

UNIX_System_Paging_Critical situation
Monitors if the virtual memory manager is working too hard to find free pages.

The formula for this situation is as follows:

```
IF System.Page_Scan_Rate GT 500
```

UNIX_User_CPU_Critical situation
Monitors if user CPU usage is system dominant and impacts users.

The formula for this situation is as follows:

```
IF System.User_CPU GE 0 AND System.User_CPU  LT 70 AND System.System_CPU GT
40
```

UNIX_System_Virtual_Memory_Warning situation
Monitors if the available virtual memory is running low.

The formula for this situation is as follows:

```
IF System.Virtual_Memory_Percent_Used GT  90
```

UNIX_WPAR_Admin_Op_Info situation
Monitors if the WPAR is running an administrative operation.

The formula for this situation is as follows:

```
AIX_WPAR_Information.Admin_Operation NE 'none'
```

UNIX_WPAR_Broken_Warning situation
Monitors if the current state of the WPAR is broken.

The formula for this situation is as follows:

```
AIX_WPAR_Information.State EQ 'Broken'
```

UNIX_WPAR_CPU_Usage_Warning situation
Monitors if the WPAR CPU usage has exceeded 95%.

The formula for this situation is as follows:

```
AIX_WPAR_CPU.WPAR_CPU_Consumed_Pct GT 95
```

UNIX_WPAR_Mem_Usage_Warning situation

Monitors if the WPAR physical memory usage has exceeded 95%.

The formula for this situation is as follows:

```
AIX_WPAR_Physical_Memory.Used_Memory_Pct GT 95
```

UNIX_WPAR_Min_CPU_Limit_Info situation

Monitors if theWPAR CPU minimum resource control is not set.

The formula for this situation is as follows:

```
AIX_WPAR_Information.RC_CPU_Limits_Min EQ 0
```

UNIX_WPAR_Min_Mem_Limit_Info situation

Monitors if the WPAR memory limit minimum resource control is set to zero.

The formula for this situation is as follows:

```
AIX_WPAR_Information.RC_Memory_Limits_Min EQ 0
```

UNIX_WPAR_RC_Inactive_Info situation

Monitors if the resource controls are not active for this WPAR.

The formula for this situation is as follows:

```
AIX_WPAR_Information.RC_Is_Active EQ 'No'
```

UNIX_WPAR_Unlim_CPU_Shares_Info situation

Monitors if the WPAR CPU resource controls are set to unlimited.

The formula for this situation is as follows:

```
AIX_WPAR_Information.RC_CPU_Shares EQ 'Unlimited' AND
AIX_WPAR_Information.RC_CPU_Limits_Hard_Max EQ 100
```

UNIX_WPAR_Unlim_Mem_Shares_Info situation

Monitors if the WPAR memory share resource controls are set to unlimited.

The formula for this situation is as follows:

```
AIX_WPAR_Information.RC_Memory_Shares EQ 'Unlimited' AND
AIX_WPAR_Information.RC_Memory_Limits_Hard_Max EQ 100
```

Chapter 6. Take Action commands reference

This chapter contains an overview of Take Action commands, references for detailed information about Take Action commands, and a description of the Take Actions command included in this monitoring agent.

About Take Action commands

Take Action commands can be run from the desktop or included in a situation or a policy.

When included in a situation, the command runs when the situation becomes true. A Take Action command in a situation is also referred to as reflex automation. When you enable a Take Action command in a situation, you automate a response to system conditions. For example, you can use a Take Action command to send a command to restart a process on the managed system or to send a text message to a cell phone.

Advanced automation uses policies to perform actions, schedule work, and automate manual tasks. A policy comprises a series of automated steps called activities that are connected to create a workflow. After an activity is completed, Tivoli Enterprise Portal receives return code feedback, and advanced automation logic responds with subsequent activities prescribed by the feedback.

More information about Take Action commands

For more information about working with Take Action commands, see the *IBM Tivoli Monitoring User's Guide*.

Predefined Take Action commands

This monitoring agent contains the following Take Action commands:

- AMS Recycle Agent Instance
- AMS Reset Agent Daily Restart Count
- AMS Start Agent
- AMS Start Agent Instance
- AMS Stop Agent
- AMS Start Management
- AMS Stop Management
- Sample_kill_Process

The remaining section of this chapter contains a description of the Take Action command. The following information is provided about Take Action commands:

Description
Which actions the command performs on the system to which it is sent

Arguments
List of arguments, if any, for the Take Action with a short description and default value for each one

Destination systems
Where the command is to be run: on the Managed System (monitoring agent) where the agent resides or on the Managing System (Tivoli Enterprise Monitoring Server) to which it is connected

Usage notes
Additional relevant notes for using the Take Actions

AMS Recycle Agent Instance

Description
Use this action to stop and start any agent with a single request. This recycle does not increase the restart count of an agent.

Arguments
Agent Name
The name of the agent as it is displayed in the Agents' Runtime Status View's Agent Name column.

Process Name
The name of the agent's process as it is displayed in the Agents' Runtime Status View's Process Name column.

Instance Name
If it exists, the name of an agent instance as it is displayed in the Agents' Runtime Status View's Instance Name column.

Process ID
The process ID of the agent process as it appears in the Agents' Runtime Status View's Process ID column.

Destination systems
Managed system

Usage notes
Not available to previous versions of the OS agents. To use this action against the OS agent, the Agent Management Services watchdog must be running.

AMS Reset Agent Daily Restart Count

Description
Use this action to reset the daily restart count of an agent to 0.

Arguments
Agent Name
The name of the agent as it is displayed in the Agents' Runtime Status View's Agent Name column.

Process Name
The name of the agent's process as it is displayed in the Agents' Runtime Status View's Process Name column.

Instance Name
If it exists, the name of an agent instance as it is displayed in the Agents' Runtime Status View's Instance Name column.

Destination systems
Managed system

Usage notes

Not available to previous versions of the OS agents. To use this action against the OS agent, the Agent Management Services watchdog must be running.

AMS Start Agent action

Description

Use this action to start an agent that is under the management of Agent Management Services. The action includes an optional input field for resetting the Daily Restart Count back to 0. This action is helpful when an agent has exceeded its maxRestartCount for the day.

Arguments

Agent Name
The name of the agent as it is displayed in the Agents' Runtime Status View's Agent Name column.

Daily Restart Count
Value indicating whether to reset the daily restart count. The value 1 indicates True, and the value 0 (default) indicates False.

Process Name
The name of the process representing the agent instance as it is displayed in the Agents' Runtime Status View's Process Name column.

Destination systems

Managed system

Usage notes

You cannot target the Monitoring Agent for UNIX OS with this action. Only the other agents being managed by Agent Management Services running on the Monitoring Agent for UNIX OS can be targeted with this action.

AMS Start Agent Instance action

Description

Use this action to start a monitoring agent instance of type ITM Windows or ITM UNIX that is under the management of Agent Management Services. The action includes an optional input field for resetting the Daily Restart Count back to 0. This action is helpful when an agent instance has exceeded its maxRestartCount for the day.

Arguments

Agent Name
The name of the agent as it is displayed in the Agents' Runtime Status View's Agent Name column.

Daily Restart Count
Value indicating whether to reset the daily restart count. The value 1 indicates True, and the value 0 (default) indicates False.

Process Name
The name of the process representing the agent instance as it is displayed in the Agents' Runtime Status View's Process Name column.

Instance Name
The name of the monitoring agent instance as it is displayed in the Agents' Runtime Status View's Instance Name column.

Destination systems
Managed system

Usage notes
You cannot target the Monitoring Agent for UNIX OS with this action. Only the other agents being managed by Agent Management Services running on the Monitoring Agent for UNIX OS can be targeted with this action.

AMS Stop Agent action

Description
Use this action to stop an agent that is under the management of Agent Management Services. The action will put a running instance of an agent into the 'Manually Stopped' state, meaning that Agent Management Services will not perform any auto-restarts. To prompt Agent Management Services to commence auto-restarting, use the AMS Start Agent command or the AMS Start Agent Instance command to manually put the agent back into a Running state.

Arguments
Process ID
By default, this argument is populated with the Process ID of the particular agent instance selected from the Tivoli Enterprise Portal. To stop all instances of an agent, such as by using the tacmd executeaction AMS Stop Agent command, leave this argument blank.

Destination systems
Managed system

Usage notes
You cannot target the Monitoring Agent for UNIX OS with this action. Only the other agents being managed by Agent Management Services running on the Monitoring Agent for UNIX OS can be targeted with this action.

AMS Start Management action

Description
Use this action to put an agent under the management of Agent Management Services. This management provides the auto-restart capability.

Destination systems
Managed system

Usage notes
You now can target the Monitoring Agent for UNIX OS with this command. Starting management of the OS Agent restarts the physical watchdog and rearms Agent Management Services. Watch of managed agents resumes. There is no change to non-OS agent management operations.

AMS Stop Management action

Description

Use this action to remove an agent from management by Agent Management Services. The action will cause the Agent Management Services watchdog to stop performing health checks and auto restarts.

Destination systems

Managed system

Usage notes

You now can target the Monitoring Agent for UNIX OS with this command. However, stopping management stops the physical watchdog and disarms Agent Management Services, which also stops watching and restarting of any managed agents. While the OS Agent is unmanaged, the AMS Start Management action will not be allowed against any other non-OS agent. The UNIX_AMS_Alert_Critical situation is activated if this take action is run on the OS agent.

Sample_kill_Process action

Description

Kills the process named in the parameter supplied and enables you to issue ad-hoc commands from the Tivoli Enterprise Portal that the Monitoring Agent for UNIX OS will run on your behalf.

Arguments

Process ID

The Process ID (PID) of the process you would like to kill.

Destination systems

Managed system

Usage notes

The kill command is run directly by the remote Monitoring Agent for UNIX OS. Because it is easy to kill processes unintentionally, you need to exercise caution if the monitoring agent is run as superuser (root).

Chapter 7. Policies reference

This chapter contains an overview of policies, references for detailed information about policies, and descriptions of the predefined policies included in this monitoring agent.

About policies

Policies are an advanced automation technique for implementing more complex workflow strategies than you can create through simple automation.

A *policy* is a set of automated system processes that can perform actions, schedule work for users, or automate manual tasks. You use the Workflow Editor to design policies. You control the order in which the policy runs a series of automated steps, which are also called activities. Policies are connected to create a workflow. After an activity is completed, Tivoli Enterprise Portal receives return code feedback and advanced automation logic responds with subsequent activities prescribed by the feedback.

Note: The predefined policies provided with this monitoring agent are not read-only. Do not edit these policies and save over them. Software updates will write over any of the changes that you make to these policies. Instead, clone the policies that you want to change to suit your enterprise.

More information about policies

For more information about working with policies, see the *IBM Tivoli Monitoring User's Guide*.

For information about using the Workflow Editor, see the *IBM Tivoli Monitoring Administrator's Guide* or the Tivoli Enterprise Portal online help.

For a list of the policies for this monitoring agent and a description of each policy, refer to the Predefined policies section below and the information in that section for each individual policy.

Predefined policies

This monitoring agent contains the following predefined policies:

- UNIX_CPU_Busy
- UNIX_Disk_Space_Full
- UNIX_Virtual_Memory_High

The remaining sections of this chapter contain descriptions of these policies, which are listed alphabetically.

UNIX_CPU_Busy policy

When the Runaway_Process and CPU_Critical situations are both true, you can choose to send a message or to terminate the runaway process (after confirmation from an administrator, if possible).

- If the termination fails, the administrator is informed, and the policy completes.

- If the termination succeeds, the policy waits and re-evaluates the CPU_Critical situation.
- If the CPU_Critical situation is still true, the administrator is informed.

UNIX_Disk_Space_Full policy

When the Disk_Space_Warning and the scratch-tmp_Disk_Full situations are both true, you can choose to perform the following actions:

- Compress all files that reside at mount point `/scratch` or `/tmp`.
- Remove all files which reside at mount point `/scratch` or `/tmp`.
- After a timeout with no user choice, echo a message.

UNIX_Virtual_Memory_High policy

When the Virtual_Memory_Warning and the Process Memory Leak situations are both true, the process identified in the Process Memory Leak event is terminated.

- If the termination fails, the administrator is informed, and the policy completes.
- If the termination succeeds, the policy waits and reevaluates the Virtual_Memory_Warning situation.
- If the Virtual_Memory_Warning situation is still true, the administrator is informed.

Chapter 8. Tivoli Common Reporting for the monitoring agent

This chapter contains a description of the data model for the Monitoring Agents for Windows OS, Linux OS, and UNIX OS reports and descriptions of these reports.

See the following additional information about using reports with this monitoring agent:

- The "Tivoli Common Reporting" chapter in the *IBM Tivoli Monitoring Administrator's Guide, V6.2.3* or later contains information about prerequisites and importing and running the reports.
- To enable Tivoli Common Reporting for monitoring agents, use the Report Installer. When requested by the Report Installer, choose the "IBM Tivoli Monitoring OS Agents Reports" package.

Complete documentation for the Tivoli Common Reporting tool is located at http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr_cog.doc/tcr_welcome.html

You must have the following prerequisite actions completed to use this function:

- IBM Tivoli Monitoring v6.2.3 must be installed with the OS agents up and running.
- The Summarization and Pruning agent must be started with or without shifts enabled.
- Historical collection must be turned on and collection started.
- Summarized tables and views must be created and populated in the Tivoli Data Warehouse.

The data model for the Monitoring Agents for Windows OS, Linux OS, and UNIX OS reports will have the following features:

- The model covers data from OS Agents only.
- The model contains a subset of the attribute groups collected by OS Agents in aggregated form for time dimension: AVG, MIN, MAX, SUM, LAT, TOT, HI, LOW. See Table 3 on page 195 for the list of included tables.
- The model includes a Managed System dimension with the Agent Type attribute (on Windows, Linux, and UNIX systems). It is placed in the IBM Tivoli Monitoring Shared Dimensions namespace.
- The model includes all the aggregations handled by Summarization and Pruning from daily to yearly including the raw data.
- The model contains forecasting based on the linear trend for the following metrics for each time dimension:

  For Linux:

  – KLZ_CPU_FCAST_XX.AVG_Idle_CPU
  – KLZ_Disk_FCAST_XX.AVG_Disk_Used_Percent
  – KLZ_VM_Stats_FCAST_XX.AVG_Used_Virtual_Storage_Pct
  – KLZ_Network_FCAST_XX.AVG_Bytes_Transmitted_per_sec
  – KLZ_Network_FCAST_XX.AVG_Bytes_Received_per_sec

  For UNIX:

  - System_FCAST_XX.AVG_Idle_CPU
  - Disk_FCAST_XX.AVG_Space_Used_Percent
  - Unix_Memory_FCAST_XX.AVG_Virtual_Storage_Pct_Used
  - Network_FCAST_XX.AVG_Transmitted_MB_Total
  - Network_FCAST_XX.AVG_Received_MB_Total

  For Windows:
  - NT_System_FCAST_XX.AVG_%_Total_Processor_Time
  - NT_Logical_Disk_FCAST_XX.AVG_%_Used
  - NT_Memory_64_FCAST_XX.AVG_Available_Usage_Percentage
  - NT_Server_FCAST_XX.AVG_Bytes_Transmitted/sec
  - NT_Server_FCAST_XX.AVG_Bytes_Received/sec
- The metrics are organized in the following way:
  - Key Metrics
    - Performance
    - Availability
  - Extended metrics
- The metric's data items names reflect the catalog attributes names with the following suffixes:
  - SUM_ into (Sum)
  - LAT_ into (Latest)
  - MIN_ into (Minimum)
  - MAX_ into (Maximum)
  - TOT_ into (Total)
  - AVG_ into (Average)
  - HI_ into (Higher)
  - LOW_ into (Lower)
- Support for raw data is provided.
- The Summarization and Pruning configuration is shown in a specific query subject (Summarization and Pruning Configuration). The result is one row that represents the most recent entry in the KSY_SUMMARIZATION_CONFIG_DV view. The query subject contains the following query items:
  - Shift Enabled. The value is 1 if the shifts hours were specified, otherwise, the value is 0.
  - Vacations Enabled. The value is 1 if the vacations days were specified, otherwise, the value is 0.
  - Peak Hours per Day. The value contains the number of peak hours specified in the shifts hours settings.
- An availability daily data query subject for each agent type is provided. Metrics are computed using the following specific availability attributes: KLZ_System_Statistics.TOT_System_Uptime, System_DV.TOT_Up_Time, NT_System.TOT_System_Up_Time_64. The calculated query items have the following meaning:
  - % Up Time. The percentage the system is available in the day.
  - % Down Time. The percentage the system is not available in the day.
  - Up Days. The portion of the day the system is available.
  - Down Days. The portion of the day the system is not available.
  - MTBSI. Mean Time Before System Interruption (in hours).

– MTTR. Mean Time To Recovery (in hours).

The following paragraphs describe the reports. In particular, they contain the required views for each one. If these views are not present, the report might not work. To ensure that the required views are present, run the following query against the Tivoli Data Warehouse:

```
DB2: select distinct "VIEWNAME" from SYSCAT.VIEWS where
"VIEWNAME" like '%V'
Oracle: select distinct "VIEW_NAME" from USER_VIEWS where
"VIEW_NAME" like '%V'
MS SQL Server: select distinct "NAME" from SYS.VIEWS where
"NAME" like '%V'
```

The following databases are supported: DB2, Oracle, and SQL Server.

The following reports are available:

- Utilization Details for Single Resource

  This report shows CPU, memory, disk, network utilization and top 10 CPU utilizing processes for a system during the selected time period in a line chart. Statistical process information is shown in all line charts (including average, upper and lower control limits). A linear trending feature is also provided and it is based on the selected forecast period.
- Utilization Details for Multiple Resources

  This report shows CPU, memory, disk and network utilization for multiple systems during the selected time period in an overlaid line chart. A linear trending feature is also provided, and it is based on the selected forecast period.
- Utilization Comparison for Single Resource

  This report shows the comparison between CPU, disk and memory utilization for a particular server, over a period of time, in an overlaid line chart.
- Utilization Comparison for Multiple Resource

  This report shows the comparison between CPU, disk and memory utilization for the selected servers over a period of time.
- Utilization Heat Chart for Single Resource

  This report helps identify patterns of utilization of a particular system over a period of time. The first column shows dates during the selected time period and the other columns represent hours during the day. The chart can be used for showing a heat chart for CPU, Memory and Disk or all three in the same report. The dates have hyperlinks that you can use to drill down to Utilization Details for Single Resource. A linear trending feature is also provided, which is based on the selected forecast period.
- Memory Utilization for Single Resource

  This report shows memory usage details for a specific system. It uses a line chart to show the percentage of virtual, physical and swap memory usage. It also provides finer memory metrics in a table.
- Memory Utilization for Multiple Resources Comparison

  This report shows memory usage details for multiple systems over a period of time. It uses three overlaid line charts for virtual, physical and swap memory.
- Top Resources Utilization

  This report shows top resources by CPU, Disk and Memory utilization. The stacked bars show average CPU used and free (in percent) for each system over the selected report period. If the number of systems is less than 20, then a bar is shown in each row. For example, there are 20 rows in the table with charts for each system. If the number of systems is more than 20, then there is a bar chart

on top with the top 20 systems, and the rest of the data is in the table. This is done to eliminate overcrowding of the bars in the chart.

- Top Situations by Status

  This report shows the top 10 situations sorted by the selected status in a bar chart, along with finer details on all the top situations, listed in a table.
- Enterprise Daily Utilization Heat Chart

  This report shows CPU, disk and memory patterns for all servers, for a select operating system type, and on a particular date. The first column lists the server names. The rest of the columns show utilization data during the day hours and the last column shows the average for the server on the selected date. You can choose to see either CPU, disk, memory or all metrics.
- Enterprise Resources List

  This report lists all the Windows, Linux and UNIX resources in the environment. By clicking on a resource name, you can drill through to see the utilization details for that resource over the last 30 days.
- Enterprise Summary

  This report shows the overall availability and utilization of all Windows, Linux and UNIX monitoring agents.
- Top Resources by Availability

  This report displays availability of the top N systems based on System Uptime over a period of time.
- Top Resources Utilization Summary Heat Chart

  This report shows top resources by CPU, Disk or Memory utilization in a summary heat chart.
- Resource Availability Comparison

  This report shows availability comparison between two or more servers.
- Top Resources by Availability (MTTR/MTBSI)

  This report displays availability trending of the top N systems based on the Mean Time Before System Interruption (MTBSI) and Mean Time To Recovery (MTTR).
- Availability Heat Chart for Single Resource

  This report helps identify patterns of resource availability over a period of time.
- CPU Utilization Comparison for Multiple Resources

  This report shows CPU usage details for multiple systems.
- CPU Utilization for Single Resource

  This report shows CPU usage details for a specific system.
- Disk Utilization for Single Resource

  This report shows the percentage of space usage for the logical disks of a particular server, over a period of time, in an overlaid line chart, along with a table that shows finer details on logical disks usage.
- Disk Utilization Comparison for Multiple Resources

  This report shows disk usage details for multiple systems, over a period of time, in two overlaid line charts.
- Situations History

  This report shows the distribution of situation events status in a pie chart, along with more detailed information on the history of situation events listed in a table.

These reports use the following attribute groups:

- **Windows agent:**
  - Logical_Disk
  - Memory
  - Process
  - Server
  - System
- **Linux agent:**
  - CPU
  - Disk
  - Network
  - Process
  - VM_Stats
- **UNIX agent:**
  - Disk
  - Network
  - Process
  - System
  - Memory
- KSY SUMMARIZATION CONFIG

The next sections in this chapter contain descriptions of the reports. For each report, the following information is included:

- Name
- Description
- Purpose
- Parameters
- Tables or views used
- Output
- Usage

One of the parameters, summarization type, has the following maximum forecast periods:

- Hourly: 60 hours in the future
- Daily: 60 days in the future
- Weekly: 1 year in the future
- Monthly: 5 years in the future
- Quarterly: no limit
- Yearly: no limit

Utilization Details for Single Resource report

Name	Utilization Details for Single Resource

Description	This report shows resources utilization for a selected server: CPU utilization, disk utilization, memory utilization, network utilization. Each metric is shown on a separate chart where data for the server is overlaid. For disk utilization, only this average value for all logical disks is shown. For network utilization, total value for all network interfaces is shown. The time frame for report data can be determined in the standard way by using the *Duration* and *Include shift periods* parameters. The server can be selected from a list of available servers by using the *OS Type* and *Servers* parameters. The forecasts can also be shown for the given period. If set, all the charts show data that ends at that date, and missing samples are determined based on linear trend computed over historical data. The report also shows the top 10 CPU utilizing processes for the selected server.
Purpose	Helps identify system performance problems related to over-utilization of key system resources. Helps identify which systems are performing poorly due to low physical memory, causing excessive paging, performing poorly due to CPU intensive tasks, or performing poorly due to other factors such as poor load balancing of applications across available systems.

Parameters	**OS Type** Determines the type of agent to work on, and is selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from or to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and is selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Servers** The server or system names for the selected OS Type are displayed in a drop-down list sorted alphabetically. You can see up to 30 system names. For more than 30 names, type the name to see the filtered list. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Forecast Period** If forecast is enabled, specifies the forecast period. **Forecast** Specifies whether forecast is enabled using a drop-down list. The list contains the following options: • Use forecast • Do not use the forecast **Show Data** Specifies if the chart data source should be displayed in a table or not.

Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **CPU Utilization:** **Windows agent:** NT_System_HV, NT_System_DV, NT_System_WV, NT_System_MV, NT_System_QV, NT_System_YV **Linux agent:** KLZ_CPU_HV, KLZ_CPU_DV, KLZ_CPU_WV, KLZ_CPU_MV, KLZ_CPU_QV, KLZ_CPU_YV **UNIX agent:** System_HV, System_DV, System_WV, System_MV, System_QV, System_YV **Disk Utilization:** **Windows agent:** NT_Logical_Disk_HV, NT_Logical_Disk_DV, NT_Logical_Disk_WV, NT_Logical_Disk_MV, NT_Logical_Disk_QV, NT_Logical_Disk_YV **Linux agent:** KLZ_Disk_HV, KLZ_Disk_DV, KLZ_Disk_WV, KLZ_Disk_MV, KLZ_Disk_QV, KLZ_Disk_YVV **UNIX agent:** Disk_HV, Disk_DV, Disk_WV, Disk_MV, Disk_QV, Disk_YV **Memory Utilization:** **Windows agent:** NT_Memory_64_HV, NT_Memory_64_DV, NT_Memory_64_WV, NT_Memory_64_MV, NT_Memory_64_QV, NT_Memory_64_YV **Linux agent:** KLZ_VM_Stats_HV, KLZ_VM_Stats_DV, KLZ_VM_Stats_WV, KLZ_VM_Stats_MV, KLZ_VM_Stats_QV, KLZ_VM_Stats_YV **UNIX agent:** Unix_Memory_HV, Unix_Memory_DV, Unix_Memory_WV, Unix_Memory_MV, Unix_Memory_QV, Unix_Memory_YV **Network Utilization:** **Windows agent:** NT_Server_HV, NT_Server_DV, NT_Server_WV, NT_Server_MV, NT_Server_QV, NT_Server_YV **Linux agent:** KLZ_Network_HV, KLZ_Network_DV, KLZ_Network_WV, KLZ_Network_MV, KLZ_Network_QV, KLZ_Network_YV **UNIX agent:** Network_HV, NetworkDV, Network_WV, Network_MV,Network_QV, Network_YV **Processes:** **Windows agent:** NT_Process_64_HV, NT_Process_64_DV, NT_Process_64_WV, NT_Process_64_MV, NT_Process_64_QV, NT_Process_64_YV **Linux agent:** KLZ_Process_HV, KLZ_Process_DV, KLZ_Process_WV, KLZ_Process_MV, KLZ_Process_QV, KLZ_Process_YV **UNIX agent:** Process_HV, Process_DV, Process_WV, Process_MV, Process_QV, Process_YV
Output	Four line charts to show CPU, disk, memory and network usage for the selected system. Each chart has 3 lines representing average, maximum and minimum % processor time used by a server over a period along with SPC data like average, upper control limit and lower control limit. A table representing the top 10 CPU utilizing processes for the selected server .

Usage	The IT administrator or manager responsible for meeting service levels based on server performance needs to receive periodic reports showing which servers are at risk of violating Service Level Agreements (SLAs) and at what times are they at most risk of violation. The same report can be used for hourly, daily, weekly, monthly, quarterly, and yearly. The ability to compare all four metrics in one chart is useful.
Drill through	On memory section title to Memory Utilization for Single Resource.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG_DV		X				
Linux	Linux CPU	KLZ_CPU	X	X	X	X	X	X
	Linux Disk	KLZ_Disk	X	X	X	X	X	X
	Linux VM Stats	KLZ_VM_Stats	X	X	X	X	X	X
	Linux Network	KLZ_Network	X	X	X	X	X	X
	Linux Process	KLZ_Process	X	X	X	X	X	X
UNIX	System	System	X	X	X	X	X	X
	Disk	Disk	X	X	X	X	X	X
	UNIX Memory	UNIX_Memory	X	X	X	X	X	X
	Network	Network	X	X	X	X	X	X
	Process	Process	X	X	X	X	X	X
Windows	System	NT_System	X	X	X	X	X	X
	Logical Disk	NT_Logical_Disk	X	X	X	X	X	X
	Memory	NT_Memory_64	X	X	X	X	X	X
	Server	NT_Server	X	X	X	X	X	X
	Process	NT_Process_64	X	X	X	X	X	X

Utilization Details for Multiple Resources report

Name	Utilization Details for Multiple Resources
Description	This report shows resources utilization for selected servers: CPU utilization, disk utilization, memory utilization, network utilization. Each metric is shown on a separate line chart where data for all servers is overlaid. For disk utilization, only average value for all logical disks is shown. For network utilization, total value for all network interfaces is shown. The time frame for report data can be determined in standard way by using the *Duration* and *Include shift periods* parameters. The servers can be selected from a list of available servers using the *OS Type* and *Servers* parameters. The forecasts can also be shown for the given period. If set, all the charts show data that ends at that date, and missing samples are determined based on linear trend computed over historical data.

Purpose	Helps identify and compare system performance problems related to over-utilization of key system resources. Helps identify which systems are performing poorly due to low physical memory, causing excessive paging, performing poorly due to CPU intensive tasks, or performing poorly due to other factors such as poor load balancing of applications across available systems.
Parameters	**OS Type** Determines the type of agent to work on and can be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and should be selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Servers** The server or system names for the selected OS Type are displayed in a drop-down list sorted alphabetically. You are able to see up to 30 system names. For more than 30 names, type the name to filter the list. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days

Parameters (Cont.)	**Forecast Period** If forecast is enabled, specifies the forecast period. **Forecast** Specifies whether forecast is enabled using a drop-down list. The list contains the following options: • Use forecast • Do not use the forecast
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **CPU Utilization:** **Windows agent:** NT_System_HV, NT_System_DV, NT_System_WV, NT_System_MV, NT_System_QV, NT_System_YV **Linux agent:** KLZ_CPU_HV, KLZ_CPU_DV, KLZ_CPU_WV, KLZ_CPU_MV, KLZ_CPU_QV, KLZ_CPU_YV **UNIX agent:** System_HV, System_DV, System_WV, System_MV, System_QV, System_YV **Disk Utilization:** **Windows agent:** NT_Logical_Disk_HV, NT_Logical_Disk_DV, NT_Logical_Disk_WV, NT_Logical_Disk_MV, NT_Logical_Disk_QV, NT_Logical_Disk_YV **Linux agent:** KLZ_Disk_HV, KLZ_Disk_DV, KLZ_Disk_WV, KLZ_Disk_MV, KLZ_Disk_QV, KLZ_Disk_YV **UNIX agent:** Disk_HV, Disk_DV, Disk_WV, Disk_MV, Disk_QV, Disk_YV **Memory Utilization:** **Windows agent:** NT_Memory_64_HV, NT_Memory_64_DV, NT_Memory_64_WV, NT_Memory_64_MV, NT_Memory_64_QV, NT_Memory_64_YV **Linux agent:** KLZ_VM_Stats_HV, KLZ_VM_Stats_DV, KLZ_VM_Stats_WV, KLZ_VM_Stats_MV, KLZ_VM_Stats_QV, KLZ_VM_Stats_YV **UNIX agent:** Unix_Memory_HV, Unix_Memory_DV, Unix_Memory_WV, Unix_Memory_MV, Unix_Memory_QV, Unix_Memory_YV **Network Utilization:** **Windows agent:** NT_Server_HV, NT_Server_DV, NT_Server_WV, NT_Server_MV, NT_Server_QV, NT_Server_YV **Linux agent:** KLZ_Network_HV, KLZ_Network_DV, KLZ_Network_WV, KLZ_Network_MV, KLZ_Network_QV, KLZ_Network_YV **UNIX agent:** Network_HV, NetworkDV, Network_WV, Network_MV,Network_QV, Network_YV
Output	Three overlaid line charts for selected systems, with each line representing the different systems. The legend is interactive.

Usage	The IT administrator or manager responsible for meeting service levels based on server performance needs to receive periodic reports showing which servers are at risk of violating Service Level Agreements (SLAs). Reports indicate which systems are overutilized or underutilized relative to a collection of systems. The report can be run hourly, daily, weekly, monthly, quarterly, and yearly.
Drill through	On legends to Utilization Details for Single Resource. On the memory section title to Memory Utilization for Multiple Resources Comparison. On the CPU section title to CPU Utilization Comparison for Multiple Resources. On the disk section title to Disk Utilization Comparison for Multiple Resources.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG_DV		X				
Linux	Linux CPU	KLZ_CPU	X	X	X	X	X	X
	Linux Disk	KLZ_Disk	X	X	X	X	X	X
	Linux VM Stats	KLZ_VM_Stats	X	X	X	X	X	X
	Linux Network	KLZ_Network	X	X	X	X	X	X
UNIX	System	System	X	X	X	X	X	X
	Disk	Disk	X	X	X	X	X	X
	UNIX Memory	UNIX_Memory	X	X	X	X	X	X
	Network	Network	X	X	X	X	X	X
Windows	System	NT_System	X	X	X	X	X	X
	Logical Disk	NT_Logical_Disk	X	X	X	X	X	X
	Memory	NT_Memory_64	X	X	X	X	X	X
	Server	NT_Server	X	X	X	X	X	X

Utilization Comparison for Single Resource report

Name	Utilization Comparison for Single Resource
Description	This report shows the comparison between CPU, disk, and memory utilization for a particular server, over a period of time, in an overlaid line chart. By clicking on the chart title, you can drill-through to see the Utilization Details for Single Resource report for the same server. The time frame for report data can be determined in the standard way by using the *Duration* and *Include shift periods* parameters. The forecasts can also be shown for the given period. If set, all the charts show data that ends at that date, and missing samples are determined based on linear trend computed over historical data.
Purpose	This report helps to compare the CPU, disk, and memory utilization of a single server.

Parameters	**OS Type** Determines the type of agent to work on, and is selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from or to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and is selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Servers** The server or system names for the selected OS Type are displayed in a drop-down list sorted alphabetically. You can see up to 30 system names. For more than 30 names, type the name to see the filtered list. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Forecast Period** If forecast is enabled, specifies the forecast period. **Forecast** Specifies whether forecast is enabled using a drop-down list. The list contains the following options: • Use forecast • Do not use the forecast

Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System_HV, NT_System_DV, NT_System_WV, NT_System_MV, NT_System_QV, NT_System_YV, NT_Logical_Disk_HV, NT_Logical_Disk_DV, NT_Logical_Disk_WV, NT_Logical_Disk_MV, NT_Logical_Disk_QV, NT_Logical_Disk_YV, NT_Memory_64_HV, NT_Memory_64_DV, NT_Memory_64_WV, NT_Memory_64_MV, NT_Memory_64_QV, NT_Memory_64_YV **Linux agent:** KLZ_CPU_HV, KLZ_CPU_DV, KLZ_CPU_WV, KLZ_CPU_MV, KLZ_CPU_QV, KLZ_CPU_YV, KLZ_Disk_HV, KLZ_Disk_DV, KLZ_Disk_WV, KLZ_Disk_MV, KLZ_Disk_QV, KLZ_Disk_YV, KLZ_VM_Stats_HV, KLZ_VM_Stats_DV, KLZ_VM_Stats_WV, KLZ_VM_Stats_MV, KLZ_VM_Stats_QV, KLZ_VM_Stats_YV **UNIX agent:** System_HV, System_DV, System_WV, System_MV, System_QV, System_YV, Disk_HV, Disk_DV, Disk_WV, Disk_MV, Disk_QV, Disk_YV, Unix_Memory_HV, Unix_Memory_DV, Unix_Memory_WV, Unix_Memory_MV, Unix_Memory_QV, Unix_Memory_YV
Output	An overlaid line chart showing the comparison between CPU, disk, and memory utilization for a particular server, over a period of time.
Usage	The IT administrator or manager responsible for meeting the server service levels needs to receive a daily report showing which servers are at risk of violating Service Level Agreements (SLAs). The report shows the overall resource utilization of a single server. The report can be run hourly, daily, weekly, monthly, quarterly, and yearly.
Drill through	By clicking on the chart title, you can drill-through to see the Utilization Details for Single Resource report for the same server.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG		X				
Linux	Linux CPU	KLZ_CPU	X	X	X	X	X	X
	Linux Disk	KLZ_Disk	X	X	X	X	X	X
	Linux VM Stats	KLZ_VM_Stats	X	X	X	X	X	X
UNIX	System	System	X	X	X	X	X	X
	Disk	Disk	X	X	X	X	X	X
	UNIX Memory	UNIX_Memory	X	X	X	X	X	X
Windows	System	NT_System	X	X	X	X	X	X
	Logical Disk	NT_Logical_Disk	X	X	X	X	X	X
	Memory	NT_Memory_64	X	X	X	X	X	X

Utilization Comparison for Multiple Resources report

Name	Utilization Comparison for Multiple Resources

Description	This report shows the comparison between CPU, disk, and memory utilization for the selected servers over a period of time. By clicking on the chart title, you can drill-through to see the corresponding Utilization Details for Multiple Resources report. By clicking on the server name, you can drill-through to see the Utilization Details for Single Resource report for the selected server. By clicking on the chart data points, you can drill-through to the corresponding CPU, Disk, or Memory Utilization for Single Resource report. The time frame for report data can be determined in the standard way by using the *Duration* and *Include shift periods* parameters. The servers can be selected from a list of available servers using the *OS Type* and *Servers* parameters. The forecasts can also be shown for the given period. If set, all the charts show data that ends at that date, and missing samples are determined based on linear trend computed over historical data.
Purpose	This report helps to compare the CPU, disk, and memory utilization for multiple servers.

Parameters	**OS Type** Determines the type of agent to work on, and is selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from or to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and is selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Servers** The server or system names for the selected OS Type are displayed in a drop-down list sorted alphabetically. You can see up to 30 system names. For more than 30 names, type the name to see the filtered list. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Forecast Period** If forecast is enabled, specifies the forecast period. **Forecast** Specifies whether forecast is enabled using a drop-down list. The list contains the following options: • Use forecast • Do not use the forecast

Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System_HV, NT_System_DV, NT_System_WV, NT_System_MV, NT_System_QV, NT_System_YV, NT_Logical_Disk_HV, NT_Logical_Disk_DV, NT_Logical_Disk_WV, NT_Logical_Disk_MV, NT_Logical_Disk_QV, NT_Logical_Disk_YV, NT_Memory_64_HV, NT_Memory_64_DV, NT_Memory_64_WV, NT_Memory_64_MV, NT_Memory_64_QV, NT_Memory_64_YV **Linux agent:** KLZ_CPU_HV, KLZ_CPU_DV, KLZ_CPU_WV, KLZ_CPU_MV, KLZ_CPU_QV, KLZ_CPU_YV, KLZ_Disk_HV, KLZ_Disk_DV, KLZ_Disk_WV, KLZ_Disk_MV, KLZ_Disk_QV, KLZ_Disk_YV, KLZ_VM_Stats_HV, KLZ_VM_Stats_DV, KLZ_VM_Stats_WV, KLZ_VM_Stats_MV, KLZ_VM_Stats_QV, KLZ_VM_Stats_YV **UNIX agent:** System_HV, System_DV, System_WV, System_MV, System_QV, System_YV, Disk_HV, Disk_DV, Disk_WV, Disk_MV, Disk_QV, Disk_YV, Unix_Memory_HV, Unix_Memory_DV, Unix_Memory_WV, Unix_Memory_MV, Unix_Memory_QV, Unix_Memory_YV
Output	Three line charts showing the CPU, disk, and memory utilization are displayed for each server selected. A table, which can be collapsed, corresponds to each chart.
Usage	The IT administrator or manager responsible for meeting the server service levels needs to receive a daily report showing which servers are at risk of violating Service Level Agreements (SLAs). The report indicates which systems are over-utilized or under-utilized relative to a collection of systems. The report can be run hourly, daily, weekly, monthly, quarterly, and yearly.
Drill through	By clicking on the chart title, you can drill-through to see the corresponding Utilization Details for Multiple Resources report. By clicking on the server name, you can drill-through to see the Utilization Details for Single Resource report for the selected server. By clicking on the chart data points, you can drill-through to the corresponding CPU, Disk or Memory Utilization for Single Resource report.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG		X				
Linux	Linux CPU	KLZ_CPU	X	X	X	X	X	X
	Linux Disk	KLZ_Disk	X	X	X	X	X	X
	Linux VM Stats	KLZ_VM_Stats	X	X	X	X	X	X
UNIX	System	System	X	X	X	X	X	X
	Disk	Disk	X	X	X	X	X	X
	UNIX Memory	UNIX_Memory	X	X	X	X	X	X
Windows	System	NT_System	X	X	X	X	X	X
	Logical Disk	NT_Logical_Disk	X	X	X	X	X	X
	Memory	NT_Memory_64	X	X	X	X	X	X

Utilization Heat Chart for Single Resource report

Name	Utilization Heat Chart for Single Resource

Description	This report helps identify patterns of utilization of a particular system over a period of time. The first column shows dates during the selected time period and the other columns represent hours during the day. The chart can be used for showing a heat chart for CPU, memory, disk or all three in the same report. The dates have hyperlinks that allow you to drill through to the Utilization Details for Single Resource report.
Purpose	Helps identify system performance of a system or server over a period of time. Shows daily patterns for utilization.

Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and should be selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Servers** The server or system names for the selected OS Type is displayed in a drop-down list sorted alphabetically. You can see up to 30 system names. For more than 30 names, type the name to filter the list. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days

Parameters (continued)	**Forecast Period** If forecast is enabled, specifies the forecast period. **Forecast** Specifies whether forecast is enabled using a drop-down list. The list contains the following options: • Use forecast • Do not use the forecast **Upper Limit for Good Status** Specifies the upper limit for good status. **Upper Limit for Fair Status** Specifies the upper limit for fair status. **Upper Limit for Warning Status** Specifies the upper limit for warning status. **Upper Limit for Bad Status and Lower Limit for Critical Status** Specifies the upper limit for bad status and the lower limit for critical status.
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **CPU Utilization:** **Windows agent:** NT_System_HV **Linux agent:** KLZ_CPU_HV **UNIX agent:** System_HV **Disk Utilization:** **Windows agent:** NT_Logical_Disk_HV **Linux agent:** KLZ_Disk_HV **UNIX agent:** Disk_HV **Memory Utilization:** **Windows agent:** NT_Memory_64_HV **Linux agent:** KLZ_VM_Stats_HV **UNIX agent:** Unix_Memory_HV
Output	A heat chart. The first column shows dates during the selected time period and the other columns represent 24 hours during the day starting with 0. The last column shows average value for that day. The report can be generated for CPU, disk or memory utilization. The timestamp is a hyperlink that you can use to drill through to a details report for CPU, disk, memory, network usage, top 10 processes for that particular system on the selected day. The thresholds for the colors can be specified in the parameters.
Usage	The IT administrator or manager responsible for meeting service levels based on server performance needs to receive periodic reports showing which servers are at risk of violating Service Level Agreements (SLAs). Reports indicate which systems are overutilized or underutilized relative to a collection of systems. The report can be run hourly, daily, weekly, monthly, quarterly, and yearly.
Drill through	On row level to Utilization Details for Single Resource.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG_DV		X				
Linux	Linux CPU	KLZ_CPU	X					
	Linux Disk	KLZ_Disk	X					
	Linux VM Stats	KLZ_VM_Stats	X					
UNIX	System	System	X					
	Disk	Disk	X					
	UNIX Memory	UNIX_Memory	X					
Windows	System	NT_System	X					
	Logical Disk	NT_Logical_Disk	X					
	Memory	NT_Memory_64	X					

Memory Utilization for Single Resource report

Name	Memory Utilization for Single Resource
Description	This report shows memory usage details for a specific system. It uses a line chart to show the percentage of virtual, physical and swap memory usage. It also provides finer memory metrics in a table. The time frame for report data can be determined in the standard way by using the *Duration* and *Include shift periods* parameters. The server can be selected from a list of available servers by using the *OS Type* and *Servers* parameters. The forecasts can also be shown for the given period. If set, all the charts show data that ends at that date, and missing samples are determined based on linear trends computed over historical data.
Purpose	Helps identify which systems are performing poorly due to low physical memory causing excessive paging.

Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and should be selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Servers** The server or system names for the selected OS Type is displayed in a drop-down list. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days

Parameters (continued)	**Forecast Period** If forecast is enabled, specifies the forecast period. **Forecast** Specifies whether forecast is enabled using a drop-down list. The list contains the following options: • Use forecast • Do not use the forecast
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_Memory_64_HV, NT_Memory_64_DV, NT_Memory_64_WV, NT_Memory_64_MV,NT_Memory_64_QV, NT_Memory_64_YV, NT_Paging_File_HV, NT_Paging_File_DV, NT_Paging_File_WV, NT_Paging_File_MV,NT_Paging_File_QV, NT_Paging_File_YV **Linux agent:** KLZ_VM_Stats_HV, KLZ_VM_Stats_DV, KLZ_VM_Stats_WV, KLZ_VM_Stats_MV,KLZ_VM_Stats_QV, KLZ_VM_Stats_YV **UNIX agent:** Unix_Memory_HV, Unix_Memory_DV, Unix_Memory_WV, Unix_Memory_MV,Unix_Memory_QV, Unix_Memory_YV
Output	A line chart showing the average usage of virtual, physical and swap memory. A table showing finer memory details.
Usage	The IT administrator or manager responsible for meeting service levels based on server performance needs to receive periodic reports showing which servers are at risk of violating Service Level Agreements (SLAs). The report indicates what is the memory health of a single system systems and if it is over-utilized or under-utilized. The report can be run hourly, daily, weekly, monthly, quarterly, and yearly.
Drill through	None.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG		X				
Linux	Linux VM Stats	KLZ_Network	X	X	X	X	X	X
	UNIX Memory	UNIX_Memory	X	X	X	X	X	X
Windows	Memory	NT_Memory_64	X	X	X	X	X	X
	Paging File	NT_Paging_File	X	X	X	X	X	X

Memory Utilization for Multiple Resources Comparison report

Name	Memory Utilization for Multiple Resources Comparison

Description	This report shows memory usage details for multiple systems over a period of time. It uses three overlaid line charts for virtual, physical and swap memory. The time frame for report data can be determined in standard way by using the *Duration* and *Include shift periods* parameters. The servers can be selected from a list of available servers by using the *OS Type* and *Servers* parameters. The forecasts can also be shown for the given period. If set, all the charts show data that ends at that date, and missing samples are determined based on linear trend computed over historical data.
Purpose	Helps identify and compare different systems behavior to identify potential memory issues due to unbalanced workload or wrong configurations. Helps identify which systems are performing poorly due to low physical memory, causing excessive paging.

Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and should be selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Servers** The server or system names for the selected OS Type is displayed in a drop-down list. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days

Parameters (continued)	**Forecast Period** If forecast is enabled, specifies the forecast period. **Forecast** Specifies whether forecast is enabled using a drop-down list. The list contains the following options: • Use forecast • Do not use the forecast
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_Memory_64_HV, NT_Memory_64_DV, NT_Memory_64_WV, NT_Memory_64_MV,NT_Memory_64_QV, NT_Memory_64_YV, NT_Paging_File_HV, NT_Paging_File_DV, NT_Paging_File_WV, NT_Paging_File_MV,NT_Paging_File_QV, NT_Paging_File_YV **Linux agent:** KLZ_VM_Stats_HV, KLZ_VM_Stats_DV, KLZ_VM_Stats_WV, KLZ_VM_Stats_MV,KLZ_VM_Stats_QV, KLZ_VM_Stats_YV **UNIX agent:** Unix_Memory_HV, Unix_Memory_DV, Unix_Memory_WV, Unix_Memory_MV,Unix_Memory_QV, Unix_Memory_YV
Output	Three overlaid line charts for selected systems, with each line representing the different systems. Each chart represents the behavior of a memory aspect.
Usage	The IT administrator or manager responsible for meeting service levels based on server performance needs to receive periodic reports showing which servers are at risk of violating Service Level Agreements (SLAs). The report indicates which systems are over-utilized or under-utilized relative to a collection of systems. The report can be run hourly, daily, weekly, monthly, quarterly, and yearly.
Drill through	On legends to Memory Utilization for Single Resource.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG		X				
Linux	Linux VM Stats	KLZ_Network	X	X	X	X	X	X
	UNIX Memory	UNIX_Memory	X	X	X	X	X	X
Windows	Memory	NT_Memory_64	X	X	X	X	X	X
	Paging File	NT_Paging_File	X	X	X	X	X	X

Top Resources Utilization report

Name	Top Resources Utilization

Description	This report shows top resources by CPU, disk and memory utilization. The stacked bars show average resource used and free (in percent) for each system over the selected report period. If the number of systems is less than 20, then a bar is shown in each row. For example, there are 20 rows in the table with charts for each system. If the number of systems is more than 20, then a bar chart is on top with the top 20 systems and the rest of the data is in the table. This is done to eliminate over-crowding of the bars in the chart.
Purpose	*CPU utilization:* Helps identify which systems are most overloaded and which have the least load based on the percentage of CPU utilization. Identifies which systems are over-utilized and which are under-utilized. *Disk utilization:* Helps identify which systems are experiencing heavy disk activity. Additionally, shows systems running low on disk space. This allows for planning the addition of hard drives or balancing of applications or data across available hard disk resources. *Memory utilization:* Helps identify growth in memory utilization which can lead to application and server outages. This allows for planning the increasing of paging space or the addition of physical memory.

Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and should be selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Number of systems** The maximum number of systems to display. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Resource** A drop-down list that you can use to choose which type of resource to display: • All • CPU • Disk • Memory

Tables or views used	**CPU utilization** • **General:** KSY_SUMMARIZATION_CONFIG_DV • **Windows agent:** NT_System_HV, NT_System_DV, NT_System_WV, NT_System_MV, NT_System_QV, NT_System_YV • **Linux agent:** KLZ_CPU_HV, KLZ_CPU_DV, KLZ_CPU_WV, KLZ_CPU_MV, KLZ_CPU_QV, KLZ_CPU_YV • **UNIX agent:** System_HV, System_DV, System_WV, System_MV, System_QV, System_YV **Disk utilization** • **Windows agent:** NT_Logical_Disk_HV, NT_Logical_Disk_DV, NT_Logical_Disk_WV, NT_Logical_Disk_MV, NT_Logical_Disk_QV, NT_Logical_Disk_YV • **Linux agent:** KLZ_Disk_HV, KLZ_Disk_DV, KLZ_Disk_WV, KLZ_Disk_MV, KLZ_Disk_QV, KLZ_Disk_YV • **UNIX agent:** Disk_HV, Disk_DV, Disk_WV, Disk_MV, Disk_QV, Disk_YV **Memory utilization** • **Windows agent:** NT_Memory_64_HV, NT_Memory_64_DV, NT_Memory_64_WV, NT_Memory_64_MV, NT_Memory_64_QV, NT_Memory_64_YV • **Linux agent:** KLZ_VM_Stats_HV, KLZ_VM_Stats_DV, KLZ_VM_Stats_WV, KLZ_VM_Stats_MV, KLZ_VM_Stats_QV, KLZ_VM_Stats_YV • **UNIX agent:** Unix_Memory_HV, Unix_Memory_DV, Unix_Memory_WV, Unix_Memory_MV, Unix_Memory_QV, Unix_Memory_YV
Output	A table is displayed with each row displaying a stacked bar representing one of the following for each system over the selected report period. • average CPU used and free (in percent) • average disk space used and free (in GB and in percent) • average memory used and free (in percent) If the number of systems is less than 20, then a bar is shown in each row. For example, there are 20 rows in the table with charts for each system. If the number of systems is more than 20, then a bar chart is on top with the top 20 systems and the rest of the data is in the table. This is done to eliminate over-crowding of the bars in the chart. The charts are interactive. By clicking on the server, the hyperlink to the Utilization Details for Single Resource is provided.
Usage	The IT administrator or manager responsible for meeting service levels based on server performance needs to receive periodic reports showing which servers are at risk of violating Service Level Agreements (SLAs). Reports indicate which systems are overutilized or underutilized relative to a collection of systems. The report can be run hourly, daily, weekly, monthly, quarterly, and yearly.
Drill through	On systems axis to Utilization Details for Single Resource.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG_DV		X				

OS Type	Attribute Group	Table	Summarization					
Linux	Linux CPU	KLZ_CPU	X	X	X	X	X	X
	Linux Disk	KLZ_Disk	X	X	X	X	X	X
	Linux VM Stats	KLZ_VM_Stats	X	X	X	X	X	X
UNIX	System	System	X	X	X	X	X	X
	Disk	Disk	X	X	X	X	X	X
	UNIX Memory	UNIX_Memory	X	X	X	X	X	X
Windows	System	NT_System	X	X	X	X	X	X
	Logical Disk	NT_Logical_Disk	X	X	X	X	X	X
	Memory	NT_Memory_64	X	X	X	X	X	X

Top Situations by Status report

Name	Top Situations by Status
Description	This report shows the top 10 situations sorted by the selected status in a bar chart, along with finer details on all the top situations, listed in a table. The time frame for the report data can be determined, in the standard way, by using the *Duration* parameter.
Purpose	Helps to analyze the top situations generating the selected event.
Parameters	**Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Status** You can choose which situation status to use in order to identify the top situations. The options are displayed in a drop-down menu where a single value can be selected between the following ones: • Acknowledged • Closed • Open • Reset • Stopped • Unknown **Aggregate Situations** You can choose if the situations should be aggregated by the Managed System and Atomize attributes or not. The default value for this parameter is Yes.
Tables or views used	**General:** CCC Logs: STATUS_HISTORY (Raw Data)

Output	A bar chart showing the top 10 situations sorted by the selected status. A table showing finer details on all the top situations sorted by the selected status.
Usage	The IT administrator or manager responsible for meeting the server service levels needs to receive periodic reports which identify the top situations generating a specific event.
Drill through	By clicking on the situation name in the table, you can drill-through to see the corresponding Situations History report.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
CCC Logs	CCC Logs	STATUS_HISTORY						

Enterprise Resources List report

Name	Enterprise Resources List
Description	This report lists all the Windows, Linux and UNIX resources in the environment. On clicking on a resource name you can drill through to see the utilization details for that resource over a period of time.
Purpose	You can use this report to see the list of OS Agents in the enterprise during a particular time.
Parameters	None
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System **Linux agent:** KLZ_CPU **UNIX agent:** System
Output	The output consists of three tables showing the resource names for Windows, Linux and UNIX. Each resource name is a hyperlink, and you can use this link to drill down to the Utilization Heat Chart for Single Resource report.
Usage	The manager responsible for meeting service levels needs to receive a weekly report of the existing systems in his environment.
Drill through	On each row in the list to Utilization Heat Chart for Single Resource.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG_DV		X				

Enterprise Daily Utilization Heat Chart report

Name	Enterprise Daily Utilization Heat Chart
Description	This report shows CPU, disk, and memory patterns for all servers, for a select operating system type, and on a particular date. The first column lists the server names. The rest of the columns show utilization data during the day hours and the last column shows the average for the server on the selected date. You can choose to see either CPU, disk, memory, or all three metrics. The date can be selected from a date prompt. The type of operating system (Linux, UNIX, Windows) can be selected from a drop down menu.
Purpose	This report helps to compare the CPU, disk and memory utilization of the machines with the same operating system in the Enterprise.
Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date** A date prompt where you can choose the date of the report. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Attribute** A drop-down list that you can use to choice what type of resource you would like to display: • All (Default) • CPU • Disk • Memory **Upper Limit for Good Status** Specifies the upper limit for good status. **Upper Limit for Fair Status** Specifies the upper limit for fair status. **Upper Limit for Warning Status** Specifies the upper limit for warning status. **Upper Limit for Bad Status and Lower Limit for Critical Status** Specifies the upper limit for bad status and the lower limit for critical status.
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System_HV, NT_Logical_Disk_HV, NT_Memory_64_HV **Linux agent:** KLZ_CPU_HV, KLZ_Disk_HV, KLZ_VM_Stats_HV **UNIX agent:** System_HV, Disk_HV, Unix_Memory_HV

Output	A heat chart per attribute (CPU, Disk, Memory) is shown for all the servers with the selected operating system. The first column lists the server names. The rest of the columns show utilization data during the day hours and the last column shows the average for the server on the selected date. You can choose to see either CPU, disk, memory or all metrics.
Usage	The IT administrator or manager responsible for meeting the server service levels needs to receive a daily report showing which servers are at risk of violating Service Level Agreements (SLAs). The report indicates which systems are over-utilized or under-utilized relative to a collection of systems.
Drill through	None.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG		X				
Linux	Linux CPU	KLZ_CPU	X					
	Linux Disk	KLZ_Disk	X					
	Linux VM Stats	KLZ_VM_Stats	X					
UNIX	System	System	X					
	Disk	Disk	X					
	UNIX Memory	UNIX_Memory	X					
Windows	System	NT_System	X					
	Logical Disk	NT_Logical_Disk	X					
	Memory	NT_Memory_64	X					

Enterprise Summary report

Name	Enterprise Summary
Description	This report shows the overall availability and utilization of all Windows, Linux and UNIX monitoring agents.
Purpose	You can use this report to compare different agent types in the environment. Note this report will run only when all 3 types of the OS agents are present in the environment.

Parameters	**Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System_DV, NT_Memory_64_DV, NT_Logical_Disk_DV **Linux agent:** KLZ_CPU_DV, KLZ_VM_Stats_DV, KLZ_Disk_DV, KLZ_System_Statistics_DV **UNIX agent:** System_DV, Disk_DV, Unix_Memory_DV
Output	The output consists of a bar chart showing a comparison of the different attributes CPU, Disk, Memory and Availability for Windows, UNIX, and Linux.
Usage	The IT administrator can see the health of the entire environment and compare the different OS types.
Drill through	On each bar to Top Resources by Utilization for the selected resource only. **Note:** This link only works for CPU, disk, and memory.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG		X				
Linux	Linux CPU	KLZ_CPU		X				
	Linux Disk	KLZ_Disk		X				
	Linux VM Stats	KLZ_VM_Stats		X				
	Linux System Statistics	KLZ_System_Statistics		X				
UNIX	System	System		X				
	Disk	Disk		X				
	UNIX Memory	UNIX_Memory		X				
Windows	System	NT_System		X				
	Logical Disk	NT_Logical_Disk		X				
	Memory	NT_Memory_64		X				

Top Resources by Availability

Name	Top Resources by Availability
Description	This report displays availability of the top N systems based on System Up time over a period of time.
Purpose	Helps identify which systems have the worst (or best) availability based on the percentage of time the system is up and running. Identifies which systems are inherently unstable.

Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Number of systems** The maximum number of systems to display. **Sort by** A drop-down list that you can use to choose how the top N list is sorted: • % Up Time • % Down Time
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System_DV **Linux agent:** KLZ_System_Statistics_DV **UNIX agent:** System_DV
Output	Stacked bar chart showing average uptime and downtime for each system over the selected report period. The bar charts are interactive and let you drill through to a heat chart for system availability.
Usage	The manager responsible for meeting service levels based on server availability needs to receive a weekly report showing which servers are at risk of violating Service Level Agreements (SLAs).

Drill through	In the bar chart to Availability Heat Chart for Single Resource.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG_DV		X				
Linux	Linux System Statistics	KLZ_System_Statistics_DV		X				
Windows	System	NT_System_DV		X				
UNIX	System	System_DV		X				

Top Resources Utilization Summary Heat Chart report

Name	Top Resources Utilization Summary Heat Chart
Description	This report shows top resources by CPU, disk, or memory utilization in a summary heat chart. By clicking on the resource name or the utilization value, you can drill through to a heat chart showing CPU, disk and memory utilization for the selected resource over the same period of time. The time frame for the report data can be determined, in the standard way, by using the *Duration* and *Include shift periods* parameters. The type of operating system (Linux, UNIX, Windows) can be selected from a drop down menu.
Purpose	This report helps to compare the top servers by CPU, disk, and memory utilization.

Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Sorting Attribute** A drop-down list that you can use to choice what type of resource you would like to display: • CPU (Default) • Disk • Memory **Number of Systems** The maximum number of servers to show in the report. **Upper Limit for Good Status** Specifies the upper limit for good status. **Upper Limit for Fair Status** Specifies the upper limit for fair status. **Upper Limit for Warning Status** Specifies the upper limit for warning status. **Upper Limit for Bad Status and Lower Limit for Critical Status** Specifies the upper limit for bad status and the lower limit for critical status.

Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System_HV, NT_Logical_Disk_HV, NT_Memory_64_HV **Linux agent:** KLZ_CPU_HV, KLZ_Disk_HV, KLZ_VM_Stats_HV **UNIX agent:** System_HV, Disk_HV, Unix_Memory_HV
Output	A heat chart with three columns for each server showing the CPU, disk, and memory utilization. The servers are sorted by CPU, disk, or memory utilization depending on the sorting attribute. The maximum number of servers shown is determined by the value of the *Number of systems* parameter.
Usage	The IT administrator or manager responsible for meeting the server service levels, needs to receive a daily report showing which servers are at risk of violating Service Level Agreements (SLAs). The report indicates which systems are over-utilized or under-utilized relative to a collection of systems.
Drill through	By clicking on the resource name or the utilization value, you can drill through to a heat chart showing CPU, disk, and memory utilization for the selected resource over the same period of time.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG		X				
Linux	Linux CPU	KLZ_CPU	X					
	Linux Disk	KLZ_Disk	X					
	Linux VM Stats	KLZ_VM_Stats	X					
UNIX	System	System	X					
	Disk	Disk	X					
	UNIX Memory	UNIX_Memory	X					
Windows	System	NT_System	X					
	Logical Disk	NT_Logical_Disk	X					
	Memory	NT_Memory_64	X					

Top Resources by Availability (MTTR/MTBSI)

Name	Top Resources by Availability (MTTR/MTBSI)
Description	This report displays availability trending of the top N systems based on the Mean Time Before System Interruption (MTBSI) and Mean Time To Recovery (MTTR).
Purpose	Help identify which systems have the worst (or best) availability based on the amount of time the system is up/running and the amount of time it takes to bring a system back online following an outage. Identifies which systems are inherently unstable.

Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Number of systems** The maximum number of systems to display. **Sort by** A drop-down list that you can use to choose how the top N list is sorted: • Mean Time To Recovery (Default) • Mean Time Before System Interruption
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System_DV **Linux agent:** KLZ_System_Statistics_DV **UNIX agent:** System_DV
Output	Stacked bar chart showing MTBSI and MTTR for each resource. An ordered table showing additional data .
Usage	The manager responsible for meeting service levels based on server availability needs to receive a weekly report showing which servers are at risk of violating Service Level Agreements (SLAs).
Drill through	None.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG_DV		X				
Linux	Linux System Statistics	KLZ_System_Statistics_DV		X				
Windows	System	NT_System_DV		X				
UNIX	System	System_DV		X				

Resource Availability Comparison

Name	Resource Availability Comparison
Description	This report shows availability comparison between two or more servers.
Purpose	Helps compare multiple systems based on availability.

Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System_DV **Linux agent:** KLZ_System_Statistics_DV **UNIX agent:** System_DV
Output	Pie charts showing % Uptime and % Downtime for selected servers. A table showing the same availability information plus details on the number of days each system is available and unavailable.
Usage	The manager responsible for meeting service levels based on server availability needs to receive a weekly report showing which servers are at risk of violating Service Level Agreements (SLAs).
Drill through	None.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG_DV		X				
Linux	Linux System Statistics	KLZ_System_Statistics_DV		X				
Windows	System	NT_System_DV		X				
UNIX	System	System_DV		X				

Availability Heat Chart for Single Resource

Name	Availability Heat Chart for Single Resource
Description	This report helps identify patterns of resource availability over a period of time.
Purpose	Helps identify system performance of a system or server over a period of time. Shows daily patterns for availability or unavailability.

Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only **Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Attribute** A drop-down list that you can use to specify which pattern to display: • % Up Time (Default) • % Down Time **Upper Limit for Good Status** Specifies the upper limit for good status. **Upper Limit for Fair Status** Specifies the upper limit for fair status. **Upper Limit for Warning Status** Specifies the upper limit for warning status. **Upper Limit for Bad Status and Lower Limit for Critical Status** Specifies the upper limit for bad status and the lower limit for critical status.
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System_HV **Linux agent:** KLZ_System_Statistics_HV **UNIX agent:** System_DV

Output	A heat chart. The first column shows dates during the selected time period and the other columns represent 24 hours during the day starting with 0. The report can also be reversed to show system downtime instead of uptime based on parameter selection. The thresholds for the colors can be specified in the parameters.
Usage	The IT administrator or manager can use this report to identify patterns of availability for a particular system over a period of time.
Drill through	None.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG_DV		X				
Linux	Linux System Statistics	KLZ_System_Statistics_HV	X					
Windows	System	NT_System_HV	X					
UNIX	System	System_DV	X					

CPU Utilization Comparison for Multiple Resources

Name	CPU Utilization Comparison for Multiple Resources
Description	This report shows CPU usage details for multiple systems, over a period of time, in three overlaid line charts for busy, user and system CPU usage on Linux and UNIX systems, and for total processor, user and privileged CPU usage on Windows systems. The time frame for the report data can be determined, in the standard way, by using the *Duration* and *Include shift period* parameters. The servers can be selected from a list of available servers using the *OS Type* and *Servers* parameters. The forecasts can also be shown for the given period. If set, all the charts show data that ends at that date, and missing samples are determined based on the linear trend computed over historical data.
Purpose	Helps to compare different system CPU usage behaviors to identify excessive CPU utilization, unbalanced workloads or wrong configurations.

Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or select from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and is selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Servers** The server or system names for the selected OS Type are displayed in a drop-down list. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only

Parameters (Continued)	**Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Forecast Period** If forecast is enabled, specifies the forecast period. **Forecast** Specifies whether forecast is enabled using a drop-down list. The list contains the following options: • Use forecast • Do not use the forecast
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System_HV, NT_System_DV, NT_System_WV, NT_System_MV, NT_System_QV,NT_System_YV **Linux agent:** KLZ_CPU_HV, KLZ_CPU_DV, KLZ_CPU_WV, KLZ_CPU_MV, KLZ_CPU_QV,KLZ_CPU_YV **UNIX agent:** System_HV, System_DV, System_WV, System_MV, System_QV, System_YV
Output	Three overlaid line charts for selected systems, with each line representing the different systems. Each chart represents the behavior of a CPU aspect. A table, which can be collapsed, corresponds to each chart.
Usage	The IT administrator or manager responsible for meeting service levels based on server performance needs to receive periodic reports showing which servers are at risk of violating Service Level Agreements (SLAs). The report indicates which systems are over-utilized or under-utilized relative to a collection of systems. The report can be run hourly, daily, weekly, monthly, quarterly, and yearly.
Drill through	On legends to CPU Utilization for Single Resource.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG		X				
Linux	Linux CPU	KLZ_CPU	X	X	X	X	X	X
Windows	System	NT_System	X	X	X	X	X	X
UNIX	System	System	X	X	X	X	X	X

CPU Utilization for Single Resource

Name	CPU Utilization for Single Resource

Description	This report shows CPU usage details for a specific system. A line chart is used to show the busy and idle CPU time trends. It also provides finer CPU metrics in a table. The time frame for the report data can be determined, in the standard way, by using the *Duration* and *Include shift period* parameters. The servers can be selected from a list of available servers using the *OS Type* and *Servers* parameters. The forecasts can also be shown for the given period. If set, all the charts show data that ends at that date, and missing samples are determined based on the linear trend computed over historical data.
Purpose	Helps identify which systems are experiencing excessive CPU usage.
Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or select from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and is selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Servers** The server or system names for the selected OS Type are displayed in a drop-down list. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only

Parameters (Continued)	**Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Forecast Period** If forecast is enabled, specifies the forecast period. **Forecast** Specifies whether forecast is enabled using a drop-down list. The list contains the following options: • Use forecast • Do not use the forecast
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_System_HV, NT_System_DV, NT_System_WV, NT_System_MV, NT_System_QV,NT_System_YV **Linux agent:** KLZ_CPU_HV, KLZ_CPU_DV, KLZ_CPU_WV, KLZ_CPU_MV, KLZ_CPU_QV,KLZ_CPU_YV **UNIX agent:** System_HV, System_DV, System_WV, System_MV, System_QV, System_YV
Output	A line chart showing busy and idle CPU time trends. A line chart showing busy and idle CPU time trends.
Usage	The IT administrator or manager responsible for meeting service levels based on server performance needs to receive periodic reports showing which servers are at risk of violating Service Level Agreements (SLAs). The report indicates what is the CPU health of a single system systems and if it is over-utilized or under-utilized. The report can be run hourly, daily, weekly, monthly, quarterly, and yearly.
Drill through	None.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG		X				
Linux	Linux CPU	KLZ_CPU	X	X	X	X	X	X
Windows	System	NT_System	X	X	X	X	X	X
UNIX	System	System	X	X	X	X	X	X

Disk Utilization for Single Resource

Name	Disk Utilization for Single Resource

Description	This report shows the percentage of space usage for the logical disks of a particular server, over a period of time, in an overlaid line chart, along with a table that shows finer details on logical disks usage. The time frame for the report data can be determined, in the standard way, by using the *Duration* and *Include shift period* parameters. The server can be selected from a list of available servers by using the *OS Type* and *Servers* parameters. The forecasts can also be shown for the given period. If set, all the charts show data that ends at that date, and missing samples are determined based on linear trend computed over historical data.
Purpose	Helps to analyze the disk utilization details of a specific machine.
Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or select from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and is selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Servers** The server or system names for the selected OS Type are displayed in a drop-down list. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only

Parameters (Continued)	**Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Include remote file systems** For Linux and UNIX systems only, it is possible to include remote file systems, such as NFS file systems, in the computation of the total average space usage percent and the total average space available in MB. **Include pseudo file systems** For Linux and UNIX systems only, it is possible to include the pseudo file systems, such as the proc file system, in the computation of the total average space usage percent and the total average space available in MB. **Forecast Period** If forecast is enabled, specifies the forecast period. **Forecast** Specifies whether forecast is enabled using a drop-down list. The list contains the following options: • Use forecast • Do not use the forecast
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_Logical_Disk_HV, NT_Logical_Disk_DV, NT_Logical_Disk_WV, NT_Logical_Disk_MV, NT_Logical_Disk_QV, NT_Logical_Disk_YV **Linux agent:** KLZ_Disk_HV, KLZ_Disk_DV, KLZ_Disk_WV, KLZ_Disk_MV, KLZ_Disk_QV, KLZ_Disk_YV **UNIX agent:** Disk_HV, Disk_DV, Disk_WV, Disk_MV, Disk_QV, Disk_YV
Output	A line chart showing the average percent space usage plotted against time. A table showing finer disk utilization details.
Usage	The IT administrator or manager responsible for meeting the server service levels, needs to receive periodic reports showing which servers are at risk of violating Service Level Agreements (SLAs). The report indicates what is the disk utilization health of a single system and which file systems are over-utilized or under-utilized. The report can be run hourly, daily, weekly, monthly, quarterly, and yearly. Note that the percent of disk usage in the report is calculated each time at run time. This approach is different from the approach used in the Tivoli Enterprise Portal Server workspace where the same metrics are instead taken directly from the % Used attribute of the Logical Disk attribute group. Due to the different units used and some rounding applied during the multiple calculations of average, the two values might vary slightly.
Drill through	None.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y

OS Type	Attribute Group	Table	Summarization					
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG		X				
Linux	Linux DISK	KLZ_DISK	X	X	X	X	X	X
Windows	Logical Disk	NT_Logical_Disk	X	X	X	X	X	X
UNIX	Disk	Disk	X	X	X	X	X	X

Disk Utilization Comparison for Multiple Resources

Name	Disk Utilization Comparison for Multiple Resources
Description	This report shows disk usage details for multiple systems, over a period of time, in two overlaid line charts. The first overlaid line chart shows the total average space usage percent plotted against time. For example, the sum of the average space usage, over a period of time, for all the file systems of a single machine, in respect to the total size of all the file systems. A linear trending feature is also provided for the total average space usage percent and it is based on the selected forecast period. The second line chart shows the total space available in megabytes plotted against time. For example, the sum of all the average space available, over a period of time, for all the file systems of a machine. By clicking on the server names in the charts legends, you can drill-through to see the corresponding Disk Utilization for Single Resource report. The time frame for the report data can be determined, in the standard way, by using the *Duration* and *Include shift period* parameters. The servers can be selected from a list of available servers using the *OS Type* and *Servers* parameters. The forecasts can also be shown for the given period. If set, all the charts show data that ends at that date, and missing samples are determined based on the linear trend computed over historical data.
Purpose	Helps to compare different file system usage behaviors to identify excessive file system utilization.

Parameters	**OS Type** Determines the type of agent to work on and should be selected from the drop-down list with the following items: • Linux • UNIX • Windows **Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or select from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Summarization Type** Determined by Summarization and Pruning and is selected from the drop-down list with the following items: • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly **Servers** The server or system names for the selected OS Type are displayed in a drop-down list. **Include shift periods** A drop-down list that you can use to select the shift periods to be included. The Peak/Off-Peak Hours period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only

Parameters (Continued)	**Include vacation periods** A drop-down list that you can use to include or exclude vacation days. The Vacation period terms refer to definitions contained in Summarization and Pruning. The list contains the following options: • All Days (Default) • Work days • Vacation days **Include remote file systems** For Linux and UNIX systems only, it is possible to include remote file systems, such as NFS file systems, in the computation of the total average space usage percent and the total average space available in MB. **Include pseudo file systems** For Linux and UNIX systems only, it is possible to the pseudo file systems, such as the proc file system, in the computation of the total average space usage percent and the total average space available in MB. **Forecast Period** If forecast is enabled, specifies the forecast period. **Forecast** Specifies whether forecast is enabled using a drop-down list. The list contains the following options: • Use forecast • Do not use the forecast
Tables or views used	**General:** KSY_SUMMARIZATION_CONFIG_DV **Windows agent:** NT_Logical_Disk_HV, NT_Logical_Disk_DV, NT_Logical_Disk_WV, NT_Logical_Disk_MV, NT_Logical_Disk_QV, NT_Logical_Disk_YV **Linux agent:** KLZ_Disk_HV, KLZ_Disk_DV, KLZ_Disk_WV, KLZ_Disk_MV, KLZ_Disk_QV, KLZ_Disk_YV **UNIX agent:** Disk_HV, Disk_DV, Disk_WV, Disk_MV, Disk_QV, Disk_YV
Output	Two overlaid line charts are shown for the selected systems, with one line for each selected system that has some historical data stored in the Tivoli Data Warehouse. Each chart represents the behavior of a different file system aspect. A table, which can be collapsed, corresponds to each chart.
Usage	The IT administrator or manager responsible for meeting the server service levels, needs to receive periodic reports showing which servers are at risk of violating Service Level Agreements (SLAs). The report indicates which systems are over-utilized or under-utilized relative to a collection of systems. The report can be run hourly, daily, weekly, monthly, quarterly, and yearly. Note that the percent of disk usage in the report is calculated each time at run time. This approach is different from the approach used in the Tivoli Enterprise Portal Server workspace where the same metrics are instead taken directly from the % Used attribute of the Logical Disk attribute group. Due to the different units used and some rounding applied during the multiple calculations of average, the two values might vary slightly.
Drill through	By clicking on one of the system names on the legends, it is possible to drill through the corresponding Disk Utilization for Single Resource report.

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
All	KSY SUMMARIZATION CONFIG	KSY_ SUMMARIZATION_ CONFIG		X				
Linux	Linux Disk	KLZ_DISK	X	X	X	X	X	X
Windows	Logical Disk	NT_Logical_Disk	X	X	X	X	X	X
UNIX	Disk	Disk	X	X	X	X	X	X

Situations History report

Name	Situations History
Description	This report shows the distribution of situation events status in a pie chart, along with more detailed information on the history of situation events listed in a table. The time frame for the report data can be determined by using Duration.
Purpose	You can use this report to analyze the history of the IBM Tivoli Monitoring situation events.

Parameters	**Date Range** Determines the range of data shown on a report. Provide the value as two border dates (from and to) or selected from the drop-down list with the following options: • All • Date Range (below) • Today • Yesterday • Last 7 days • Last 30 days • Last 90 days • Last 355 days • Current week • Current month • Current Year to Date • Last week • Last month • Last Year **Status** You can apply a filter on the situations event data set by specifying the status in a multi-select value prompt where one or multiple status value can be selected from the following: • Acknowledged • Closed • Open • Reset • Stopped • Unknown **Managed System Filter** You can apply a filter on the situations events data set by specifying a regular expression that the managed system attribute should follow. This filter can contain the two following wildcard characters: the percent sign (%), which matches zero or more characters, and the underscore sign (_), which matches a single character. The default value for the regular expression is the percent sign, and, by default, all the managed system are selected. The escape character for the underscore and percent signs is the backslash character (\). The empty string for the Situation Name Filter can be specified through two single quotation marks (' '). **Situation Name Filter** You can apply a filter on the situations events data set by specifying a regular expression that the situation name attribute should follow. This filter can contain the two following wildcard characters: the percent sign (%), which matches zero or more characters, and the underscore sign (_), which matches a single character. The default value for the regular expression is the percent sign, and, by default, all the situation names are selected. The escape character for the underscore and percent signs is the backslash character (\).
Tables or views used	**General:** CCC Logs: STATUS_HISTORY (Raw Data)
Output	A pie chart showing the distribution of situation events status. A table showing more detailed information on situation status history.

Usage	The IT administrator or manager responsible for meeting the server service levels, needs to receive periodic reports showing which is the situation event status distribution.
Drill through	None

The following table includes information about the historical collection configuration:

OS Type	Attribute Group	Table	Summarization					
			H	D	W	M	Q	Y
CCC Logs	CCC Logs	STATUS_HISTORY						

Creating custom queries and reports

You can create your own queries and reports using the models and reports that have been documented in the subsections above by completing the following steps:

1. Read the instructions for enabling historical collection found in the *Tivoli Enterprise Portal User's Guide*.
2. Check in Table 3 below for the list of the attribute groups that are supported by the data model and are found in the Tivoli Data Warehouse database.
3. Enable historical collection for these supported attribute groups and configure the summarization settings. All of the summarization settings are supported.
4. After the database is populated, use the model leveraging in Query Studio and Report Studio.

Table 3. Attributes groups supported by the data model

Agent	Attribute groups	Tables
Linux	Linux CPU Averages	KLZ_CPU_Averages
	Linux CPU	KLZ_CPU
	Linux Disk	KLZ_Disk
	Linux Network	KLZ_Network
	Linux Process	KLZ_Process
	Linux VM Stats	KLZ_VM_Stats
	Linux Disk IO	KLZ_Disk_IO
	Linux Disk Usage Trends	KLZ_Disk_Usage_Trends
	Linux IO Ext	KLZ_IO_Ext
	Linux NFS Statistics	KLZ_NFS_Statistics
	Linux Process User Info	KLZ_Process_User_Info
	Linux RPC Statistics	KLZ_RPC_Statistics
	Linux Sockets Detail	KLZ_Sockets_Detail
	Linux Sockets Status	KLZ_Sockets_Status
	Linux Swap Rate	KLZ_Swap_Rate
	Linux System Statistics	KLZ_System_Statistics
	Linux User Login	KLZ_User_Login

Table 3. Attributes groups supported by the data model (continued)

Agent	Attribute groups	Tables
UNIX	Disk	Disk
	Network	Network
	Process	Process
	Unix Memory	Unix_Memory
	System	System
	Disk Performance	Disk_Performance
	NFS and RPC Statistics	N_F_S_and_R_P_C_Statistics
	SMP CPU	SMP_CPU
	Solaris Zones	Solaris_Zones
	User	User
Windows	Logical Disk Hourly	NT_Logical_Disk
	Memory Hourly	NT_Memory_64
	Network Interface Hourly	NT_Network_Interface
	Process Hourly	NT_Process_64
	Server Hourly	NT_Server
	System Hourly	NT_Process_64
	ICMP Statistics Hourly	ICMP_Statistics
	IP Statistics Hourly	IP_Statistics
	Cache Hourly	NT_Cache
	Device Dependencies Hourly	NT_Device_Dependencies
	Devices Hourly	NT_Devices
	Event Log Hourly	NT_Event_Log
	Monitored Logs Report Hourly	NT_Monitored_Logs_Report
	Network Port Hourly	NT_Network_Port
	Objects Hourly	NT_Objects
	Paging File Hourly	NT_Paging_File
	Physical Disk Hourly	NT_Physical_Disk
	Printer Hourly	NT_Printer
	Processor Hourly	NT_Processor
	Processor Summary Hourly	NT_Processor_Summary
	Redirector Hourly	NT_Redirector
	Server Work Queues Hourly	NT_Server_Work_Queues_64
	Service Dependencies Hourly	NT_Service_Dependencies
	Services Hourly	NT_Services
	Thread Hourly	NT_Thread
	Print Queue Hourly	Print_Queue
	Process IO Hourly	Process_IO
	TCP Statistics Hourly	TCP_Statistics
	UDP Statistics Hourly	UDP_Statistics

Note: There is a subset of tables that are visible in the model, but cannot be used in custom queries and reports. These tables are contained in the following folders:

- Forecast Hourly
- Forecast Daily
- Forecast Weekly
- Forecast Monthly
- Forecast Quarterly
- Forecast Yearly

Chapter 9. Troubleshooting

This chapter explains how to troubleshoot the IBM Tivoli Monitoring: UNIX OS Agent. Troubleshooting, or problem determination, is the process of determining why a certain product is malfunctioning.

Note: You can resolve some problems by ensuring that your system matches the system requirements listed in Chapter 2, "Requirements for the monitoring agent," on page 7.

This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information. Also see "Support information" on page 222 for other problem-solving options.

Gathering product information for IBM Software Support

Before contacting IBM Software Support about a problem you are experiencing with this product, gather the following information that relates to the problem:

Table 4. Information to gather before contacting IBM Software Support

Information type	Description
Log files	Collect trace log files from failing systems. Most logs are located in a `logs` subdirectory on the host computer. See "Trace logging" on page 200 for lists of all trace log files and their locations. See the *IBM Tivoli Monitoring User's Guide* for general information about the IBM Tivoli Monitoring environment.
UNIX information	• Version number and patch level • Sample application data file (if monitoring a file)
Operating system	Operating system version number and patch level
Messages	Messages and other information displayed on the screen
Version numbers for IBM Tivoli Monitoring	Version number of the following members of the monitoring environment: • IBM Tivoli Monitoring. Also provide the patch level, if available. • IBM Tivoli Monitoring: UNIX OS Agent
Screen captures	Screen captures of incorrect output, if any.
(*UNIX only*) Core dump files	If the system stops on UNIX systems, collect core dump file from *install_dir*/bin directory, where *install_dir* is the directory path where you installed the monitoring agent.

Built-in troubleshooting features

The primary troubleshooting feature in the IBM Tivoli Monitoring: UNIX OS Agent is logging. *Logging* refers to the text messages and trace data generated by the IBM Tivoli Monitoring: UNIX OS Agent. Messages and trace data are sent to a file.

Trace data captures transient information about the current operating environment when a component or application fails to operate as designed. IBM Software Support personnel use the captured trace information to determine the source of an error or unexpected condition. See "Trace logging" on page 200 for more information.

Problem classification

The following types of problems might occur with the IBM Tivoli Monitoring: UNIX OS Agent:

- Installation and configuration
- General usage and operation
- Display of monitoring data
- Take Action commands

This chapter provides symptom descriptions and detailed workarounds for these problems, as well as describing the logging capabilities of the monitoring agent. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Trace logging

Trace logs capture information about the operating environment when component software fails to operate as intended. The principal log type is the RAS (Reliability, Availability, and Serviceability) trace log. These logs are in the English language only. The RAS trace log mechanism is available for all components of IBM Tivoli Monitoring. Most logs are located in a `logs` subdirectory on the host computer. See the following sections to learn how to configure and use trace logging:

- "Principal trace log files" on page 201
- "Examples: using trace logs" on page 203
- "Setting RAS trace parameters" on page 204

Note: The documentation refers to the RAS facility in IBM Tivoli Monitoring as "RAS1".

The default configuration for trace logging, such as whether trace logging is enabled or disabled and trace level, depends on the source of the trace logging. Trace logging is always enabled.

Typically, IBM Software Support applies specialized knowledge to analyze trace logs to determine the source of problems. However, you can open trace logs in a text editor such as **vi** to learn some basic facts about your IBM Tivoli Monitoring environment as described in "Examples: using trace logs" on page 203.

Overview of log file management

Table 5 on page 202 provides the names, locations, and descriptions of RAS1 log files. The log file names adhere to the following naming convention:

hostname_product_program_timestamp-nn.log

where:

- *hostname* is the host name of the computer on which the monitoring component is running.
- *product* is the two-character product code. For Monitoring Agent for UNIX OS, the product code is ux.
- *program* is the name of the program being run.
- *timestamp* is an 8-character hexadecimal timestamp representing the time at which the program started.

- *nn* is a rolling log suffix. See “Examples of trace logging” for details of log rolling.

Examples of trace logging

For example, if a UNIX monitoring agent is running on computer "server01", the RAS log file for the Monitoring Agent for UNIX OS might be named as follows:

```
server01_ux_kuxagent_437fc59-01.log
```

For long-running programs, the *nn* suffix is used to maintain a short history of log files for that startup of the program. For example, the kuxagent program might have a series of log files as follows:

```
server01_ux_kuxagent_437fc59-01.log
server01_ux_kuxagent_437fc59-02.log
server01_ux_kuxagent_437fc59-03.log
```

As the program runs, the first log (*nn*=01) is preserved because it contains program startup information. The remaining logs "roll." In other words, when the set of numbered logs reach a maximum size, the remaining logs are overwritten in sequence.

Each time a program is started, a new timestamp is assigned to maintain a short program history. For example, if the Monitoring Agent for UNIX OS is started twice, it might have log files as follows:

```
server01_ux_kuxagent_437fc59-01.log
server01_ux_kuxagent_437fc59-02.log
server01_ux_kuxagent_437fc59-03.log

server01_ux_kuxagent_537fc59-01.log
server01_ux_kuxagent_537fc59-02.log
server01_ux_kuxagent_537fc59-03.log
```

Each program that is started has its own log file. For example, the Monitoring Agent for UNIX OS would have agent logs in this format:

```
server01_ux_kuxagent_437fc59-01.log
```

Other logs, such as logs for UNIX collector processes and Take Action commands, have a similar syntax as in the following example:

```
server01_ux_ifstat_447fc59-01.log
```

where **ifstat** is the program name.

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

Principal trace log files

Table 5 on page 202 contains locations, file names, and descriptions of trace logs that can help determine the source of problems with agents.

Table 5. Trace log files for troubleshooting agents

System where log is located	File name and path	Description
On the computer that hosts the monitoring agent See "Definitions of variables" on page 203 for descriptions of the variables in the file names in column two.	The RAS1 log files are named *hostname*_ux_*program_timestamp-nn*.log and are located in the *install_dir*/logs path. **Note:** File names for RAS1 logs include a hexadecimal timestamp. **Also on UNIX, a log with a decimal timestamp is provided:** *hostname*_ux_*timestamp*.log and *hostname*_ux_*timestamp*.pid*nnnnn* in the *install_dir*/*logs* path, where *nnnnn* is the process ID number.	Traces activity of the monitoring agent. **Note:** Other logs, such as logs for UNIX collector processes and Take Action commands (if available), have a similar syntax and are located in this directory path.
	The *.LG0 file is located in the *install_dir*/logs path.	A new version of this file is generated every time the agent is restarted. IBM Tivoli Monitoring generates one backup copy of the *.LG0 file with the tag .LG1. View .LG0 to learn the following details regarding the current monitoring session: • Status of connectivity with the monitoring server. • Situations that were running. • The success or failure status of Take Action commands.
On the Tivoli Enterprise Monitoring Server See "Definitions of variables" on page 203 for descriptions of the variables in the file names in column two.	**On UNIX:** The candle_installation.log file in the *install_dir*/logs path. **On Windows:** The file in the *install_dir*\InstallITM path.	Provides details about products that are installed. **Note:** Trace logging is enabled by default. A configuration step is not required to enable this tracing.
	The Warehouse_Configuration.log file is located in the following path on Windows: *install_dir*\InstallITM.	Provides details about the configuration of data warehousing for historical reporting.
	The RAS1 log file is named *hostname*_ms_*timestamp-nn*.log and is located in the following path: • **On Windows:** *install_dir*\logs • **On UNIX:** *install_dir*/logs **Note:** File names for RAS1 logs include a hexadecimal timestamp **Also on UNIX, a log with a decimal timestamp is provided:** *hostname*_ms_*timestamp*.log and *hostname*_ms_*timestamp*.pid*nnnnn* in the *install_dir*/*logs* path, where *nnnnn* is the process ID number.	Traces activity on the monitoring server.

Table 5. Trace log files for troubleshooting agents (continued)

System where log is located	File name and path	Description
On the Tivoli Enterprise Portal Server See "Definitions of variables" for descriptions of the variables in the file names in column two.	The RAS1 log file is named *hostname*_cq_*timestamp-nn*.log and is located in the following path: • **On Windows:** *install_dir*\logs • **On UNIX:** *install_dir*/logs **Note:** File names for RAS1 logs include a hexadecimal timestamp **Also on UNIX, a log with a decimal timestamp is provided:** *hostname*_cq_*timestamp*.log and *hostname*_cq_*timestamp*.pid*nnnnn* in the *install_dir/logs* path, where *nnnnn* is the process ID number.	Traces activity on the portal server.
	The TEPS_ODBC.log file is located in the following path on Windows: *install_dir*\InstallITM.	When you enable historical reporting, this log file traces the status of the warehouse proxy agent.
Definitions of variables for RAS1 logs: • *hostname* is the host name of the computer on which the agent is running. • *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. • *product* is the two character product code. For Monitoring Agent for UNIX OS, the product code is ux. • *program* is the name of the program being run. • *timestamp* is an eight-character hexadecimal timestamp representing the time at which the program started. • *nn* is a rolling log suffix. See "Examples of trace logging" on page 201 for details of log rolling.		

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

Examples: using trace logs

Typically IBM Software Support applies specialized knowledge to analyze trace logs to determine the source of problems. However, you can open trace logs in a text editor such as **vi** to learn some basic facts about your IBM Tivoli Monitoring environment. You can use the **ls -ltr** command to list the log files in the *install_dir*/logs directories, sorted by time they were last updated.

Example one

This excerpt shows the typical log for a failed connection between a monitoring agent and a monitoring server with the host name **server1a**:

```
(Thursday, August 11, 2005, 08:21:30-{94C}kdcl0cl.c,105,"KDCL0_ClientLookup") status=1c020006,
    "location server unavailable", ncs/KDC1_STC_SERVER_UNAVAILABLE
(Thursday, August 11, 2005, 08:21:35-{94C}kraarreg.cpp,1157,"LookupProxy") Unable to connect to
    broker at ip.pipe:: status=0, "success", ncs/KDC1_STC_OK
(Thursday, August 11, 2005, 08:21:35-{94C}kraarreg.cpp,1402,"FindProxyUsingLocalLookup") Unable
    to find running CMS on CT_CMSLIST <IP.PIPE:#server1a>
```

Example two

The following excerpts from the trace log *for the monitoring server* show the status of an agent, identified here as "Remote node." The name of the computer where the agent is running is **SERVER5B**:

```
(42C039F9.0000-6A4:kpxreqhb.cpp,649,"HeartbeatInserter") Remote node SERVER5B:KUX is ON-LINE.
. . .
(42C3079B.0000-6A4:kpxreqhb.cpp,644,"HeartbeatInserter") Remote node SERVER5B:KUX is OFF-LINE.
```

Key points regarding the preceding excerpt:

- The monitoring server appends the **KUX** product code to the server name to form a unique name (`SERVER5B:KUX`) for this instance of Monitoring Agent for UNIX OS. This unique name enables you to distinguish multiple monitoring products that might be running on **SERVER5B**.
- The log shows when the agent started (ON-LINE) and later stopped (OFF-LINE) in the environment.
- For the sake of brevity an ellipsis (...) represents the series of trace log entries that were generated while the agent was running.
- Between the ON-LINE and OFF-LINE log entries, the agent was communicating with the monitoring server.
- The ON-LINE and OFF-LINE log entries are always available in the trace log. All trace levels that are described in "Setting RAS trace parameters" provide these entries.

Setting RAS trace parameters

Objective

Pinpoint a problem by setting detailed tracing of individual components of the monitoring agent and modules.

Background Information

Monitoring Agent for UNIX OS uses RAS1 tracing and generates the logs described in Table 5 on page 202. The default RAS1 trace level is ERROR.

RAS1 tracing has control parameters to manage to the size and number of RAS1 logs. Use the procedure described in this section to set the parameters.

Note: The **KBB_RAS1_LOG** parameter also provides for the specification of the log file directory, log file name, and the inventory control file directory and name. Do not modify these values or log information can be lost.

Before you begin

See "Overview of log file management" on page 200 to ensure that you understand log rolling and can reference the correct log files when you managing log file generation.

After you finish

Monitor the size of the **logs** directory. Default behavior can generate a total of 45 to 60 MB for each agent that is running on a computer. For example, each database instance that you monitor could generate 45 to 60 MB of log data. See the "Procedure" section to learn how to adjust file size and numbers of log files to prevent logging activity from occupying too much disk space.

Regularly prune log files other than the RAS1 log files in the **logs** directory. Unlike the RAS1 log files which are pruned automatically, other log types can grow indefinitely, for example, the logs in Table 5 on page 202 that include a process ID number (PID).

Consider using collector trace logs (described in Table 5 on page 202) as an additional source of troubleshooting information.

Note: The **KDC_DEBUG** setting and the Maximum error tracing setting can generate a large amount of trace logging. Use them only temporarily, while you are troubleshooting problems. Otherwise, the logs can occupy excessive amounts of hard disk space.

Procedure

Specify RAS1 trace options in the *install_dir*/config/ux.ini file. You can manually edit the configuration file to set trace logging:

1. Open the trace options file: */install_dir*/config/ux.ini.
2. Edit the line that begins with **KBB_RAS1=** to set trace logging preferences.

   For example, if you want detailed trace logging, set the Maximum Tracing option:

   ```
   export KBB_RAS1='ERROR (UNIT:kux ALL) (UNIT:kra ALL)'
   ```
3. Edit the line that begins with **KBB_RAS1_LOG=** to manage the generation of log files:
   - Edit the following parameters to adjust the number of rolling log files and their size.
     - **MAXFILES**: the total number of files that are to be kept for all startups of a given program. Once this value is exceeded, the oldest log files are discarded. Default value is 9.
     - **LIMIT**: the maximum size, in megabytes (MB) of a RAS1 log file. Default value is 5.
   - IBM Software Support might guide you to modify the following parameters:
     - **COUNT**: the number of log files to keep in the rolling cycle of one program startup. Default value is 3.
     - **PRESERVE**: the number of files that are not to be reused in the rolling cycle of one program startup. Default value is 1.

   Note: The **KBB_RAS1_LOG** parameter also provides for the specification of the log file directory, log file name, and the inventory control file directory and name. Do not modify these values or log information can be lost.
4. Restart the monitoring agent so that your changes take effect.

Problems and workarounds

The following sections provide symptoms and workarounds for problems that might occur with Monitoring Agent for UNIX OS:

- "Installation and configuration troubleshooting" on page 206
- "Agent troubleshooting" on page 210
- "Tivoli Enterprise Portal troubleshooting" on page 215
- "Workspace troubleshooting" on page 216
- "Troubleshooting for remote deployment" on page 216
- "Situation troubleshooting" on page 217
- "Take Action command troubleshooting" on page 221
- "Troubleshooting for UNIX" on page 221

Note: You can resolve some problems by ensuring that your system matches the system requirements listed in Chapter 2, "Requirements for the monitoring agent," on page 7.

This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Installation and configuration troubleshooting

This section provides tables that show solutions for installation, configuration, and uninstallation problems.

Agent upgrade and restart using non-root

The monitoring agent can run using a non-root user ID on UNIX and Linux systems. This can be done by running the **itmcmd agent start** command while logged in as a non-root user, and this can be done remotely by deploying the agent using the **Run As** option on the GUI or using the **_UNIX_STARTUP_.Username** option on the **tacmd addSystem** command line. If the agent is running using a non-root user ID, and then the agent is upgraded, restarted remotely, restarted as a result of a system reboot, or the **itmcmd agent start** is run using the root user ID, then the monitoring agent subsequently runs as the root user. To confirm the user ID that the monitoring agent is using, run the following command:

```
itm_install/bin/cinfo -r
```

If the agent is using root, and that is not the desired user ID, then use the following steps to restart the agent:

1. Log in as root.
2. Run the **itmcmd agent stop** command.
3. Log in (or 'su') to the user ID that you want the agent to run as.
4. Run the **itmcmd agent start** command.

If the agent was running as root because of a system reboot, then edit the startup file using the following steps so that the appropriate user ID is used the next time the system is rebooted:

1. Look at *install_dir*/registry/AutoStart, and get *NUM*.
2. Edit the autostart for your operating system:

   The location of the startup file is platform dependent as follows:

   - AIX: /etc/rc.itm*NUM*
   - HP-UX: /sbin/init.d/ITMAgents*NUM*
   - Linux: /etc/init.d/ITMAgents*NUM*
   - Solaris: /etc/init.d/ITMAgents*NUM*
3. Add entries for your operating system using the following command:

   ```
   /usr/bin/su - instancename
   -c "install_dir/bin/itmcmd agent
   -h install_dir
   -o instancename
   start product_code"
   ```

   Where:

   instancename
   : Name of the instance

   install_dir
   : Name of the directory

   product_code
   : 2-character product code for the agent, for example, ux for the Monitoring Agent for UNIX OS

Examples:

- For AIX, add entries with the following format:

```
su - USER -c " /opt/IBM/ITM/bin/itmcmd agent
-o INSTANCE start ux"
```

  Where:

  USER Name of the user

  INSTANCE
  Name of the instance

- For Linux, HP_UX, and Solaris, add entries with the following format:

```
/bin/su - USER -c " /opt/IBM/ITM/bin/itmcmd agent
-o INSTANCE start ux >/dev/null 2>&1"
```

  Where:

  USER Name of the user

  INSTANCE
  Name of the instance

4. Repeat Steps 1 through 3 for all occurrences of stop.
5. Save the file.

Table 6. Problems and solutions for installation and configuration

Problem	Solution
Installation fails on HPUX11. The log for the monitoring agent shows the message listed in the next row of this table.	You must install the **PHSS_30966** patch on the HPUX system. See the Web site listed in the next row of this table for details.

When a patch is missing on HPUX11, the following message is generated:
`/usr/lib/pa20_64/dld.sl: Unsatisfied code symbol 'dladdr' in load module \`
`'/opt/IBM/ITM/tmaitm6/hp116/lib/libkbb.sl'`
 	
The following Web site provides details about the required patch for HPUX11:
`http://www2.itrc.hp.com/service/patch/patchDetail.do?patchid=PHSS_30966&admit=-1335382922+112672773755 \`
`6+28353475`	

When you upgrade to IBM Tivoli Monitoring, you might need to apply fixpacks to Candle, Version 350, agents.	Fixpacks for Candle, Version 350, are delivered as each monitoring agent is upgraded to IBM Tivoli Monitoring. **Note:** The IBM Tivoli Monitoring download image or CD provides application fixpacks for the monitoring agents that are installed from that CD (for example, the agents for operating systems such as Windows, Linux, UNIX, and i5/OS®). The upgrade software for other agents is located on the download image or CDs for that specific monitoring agent, such as the agents for database applications. If you do not upgrade the monitoring agent to IBM Tivoli Monitoring, the agent continues to work. However, you must upgrade to have all the functionality that IBM Tivoli Monitoring offers.
Presentation files and customized OMEGAMON® screens for Candle monitoring agents need to be upgraded to a new Linux on z/Series system.	The upgrade from version 350 to IBM Tivoli Monitoring handles export of the presentation files and the customized OMEGAMON screens.

Table 6. Problems and solutions for installation and configuration (continued)

Problem	Solution
The product fails to do a monitoring activity that requires read, write, or execute permissions. For example, the product might fail to run a Take Action command or read a log.	The monitoring agent must have the permissions necessary to perform requested actions. For example, if the user ID you used to log onto the system to install the monitoring agent (locally or remotely) does not have the permission to perform a monitoring operation (such as running a command), the monitoring agent is not able perform the operation.
While installing the agent from a CD, the following message is displayed and you are not able to continue the installation: `install.sh warning: unarchive of "/cdrom/unix/cienv1.tar" may have failed`	This error is caused by low disk space. Although the `install.sh` script indicates that it is ready to install the agent software, the script considers the size of *all* tar files, not the size of all the files that are contained within the tar file.Run the **df -k** command to check whether the file systems have enough space to install agents.
The Monitoring Agent for UNIX OS repeatedly restarts.	You can collect data to analyze this problem as follows: 1. Access the *install_dir*`/config/ux.ini` file, which is described in "Setting RAS trace parameters" on page 204. 2. Add the following line: `KBB_SIG1=trace –dumpoff`
Agents in the monitoring environment use different communication protocols. For example, some agents have security enabled and others do not.	Configure both the monitoring server and the Warehouse proxy server to accept multiple protocols, as described in the *IBM Tivoli Monitoring Installation and Setup Guide*.
Creating a firewall partition file: The partition file enables an agent to connect to the monitoring server through a firewall.	**How it works:** When the agents start, they search `KDCPARTITION.TXT` for the following matches: • An entry that matches the partition name **OUTSIDE**. • An entry that also includes a valid external address. For more information, see the *IBM Tivoli Monitoring Installation and Setup Guide*.
The Monitoring Agent for UNIX OS is started and running but not displaying data in the Tivoli Enterprise Portal.	Check the following issues: 1. Check the Monitoring Agent for UNIX OS log files to see whether there are connection problems. 2. If there are no connection problems, check whether the agent has terminated. (Search for the word "terminated" in the log.) 3. If the agent is not terminated, confirm that you have added application support for the Monitoring Agent for UNIX in the Tivoli Enterprise Monitoring Server, as described in the *IBM Tivoli Monitoring Installation and Setup Guide*.
You successfully migrate an OMEGAMON monitoring agent to IBM Tivoli Monitoring, Version 6.1.0. However, when you configure historical data collection, you see an error message that includes, `Attribute name may be invalid, or attribute file not installed for warehouse agent`.	Install the agent's application support files on the Tivoli Enterprise Monitoring Server, using the following steps: 1. Open the Manage Tivoli Enterprise Monitoring Services window. 2. Right-click the name of the monitoring server. 3. Select **Advanced > Add TEMS Application Support** in the pop-up menu. Add application support if any for any agent that is missing from the list. See the *IBM Tivoli Monitoring Installation and Setup Guide* for more information on adding application support. Ensure that the agent's application support files are pushed to the system that houses the Warehouse Proxy Agent. The Warehouse Proxy must be able to access the short attribute names for tables and columns. That way, if the longer versions of these names exceed the limits of the Warehouse database, the shorter names can be substituted.

Table 7. General problems and solutions for uninstallation

Problem	Solution
The way to remove inactive managed systems (systems whose status is OFFLINE) from the Enterprise navigation tree in the portal is not obvious.	When you want to remove a managed system from the navigation tree, complete the following steps: 1. Click **Enterprise** in the navigation tree. 2. Right-click **Workspace -> Managed System Status**. 3. Right-click the offline managed system and select **Clear offline entry**.

Unique names for monitoring components

If you have multiple instances of a monitoring agent, you must decide how to name the monitoring agents. This name is intended to uniquely identify that monitoring agent. The agent's default name is composed of three qualifiers:

- Optional instance name
- Computer network host name
- Agent product node type

An agent name truncation problem can occur when the network domain name is included in the network host name portion of the agent name. For example, instead of just the host name myhost1 being used, the resulting host name might be myhost1.acme.north.prod.com. Inclusion of the network domain name causes the agent name in the example above to expand to SERVER1:myhost1.acme.north.prod.com:KXX. This resulting name is 39 characters long. It is truncated to 32 characters resulting in the name SERVER1:myhost1.acme.north.prod.

The agent name truncation is only a problem if there is more than one monitoring agent on the same system. In this case, the agent name truncation can result in collisions between agent products attempting to register using the same truncated name value. When truncated agent names collide on the same system, this can lead to Tivoli Enterprise Monitoring Server problems with corrupted EIB tables. The agent name collision in the Tivoli Enterprise Monitoring Server might cause a registered name to be associated with the wrong product.

In general, create names that are short but meaningful within your environment. Use the following guidelines:

- Each name must be unique. One name cannot match another monitoring agent name exactly.
- Each name must begin with an alpha character.
- Do not use blanks or special characters, including $, #, and @.
- Each name must be between 2 and 32 characters in length.
- Monitoring agent naming is case-sensitive on all operating systems.

Create the names by completing the following steps:

1. Open the configuration file for the monitoring agent, which is located in the following path:
   - **On Windows:** *&install_dir;*\tmaitm6\K*product_code*CMA.INI. For example, the product code for the Monitoring Agent for Windows OS is NT and the file name is KNTCMA.INI.
   - **On UNIX and Linux:** *install_dir*/tmaitm6/*product_code*.ini and *product_code*.config. For example, the file names for the Monitoring Agent for UNIX OS is ux.ini and ux.config.

2. Find the line the begins with **CTIRA_HOSTNAME=**.
3. Type a new name for host name that is a unique, shorter name for the host computer. The final concatenated name including the subsystem name, new host name, and UX, cannot be longer than 32 characters.

   Note: You must ensure that the resulting name is unique with respect to any existing monitoring component that was previously registered with the Tivoli Enterprise Monitoring Server.
4. Save the file.
5. Restart the agent.
6. If you do not find the files mentioned in Step 1, perform the workarounds listed in the next paragraph.

If you do not find the files mentioned in the preceding steps, perform the following workarounds:

1. Change **CTIRA_HOSTNAME** environment variable in the configuration file of the monitoring agent.
   - Find the KUXENV file in the same path mentioned in the preceding row.
   - For z/OS® agents, find the **RKANPAR** library.
   - For i5/OS agents, find the **QAUTOTMP/KMSPARM** library in member **KBBENV**.
2. If you cannot find the **CTIRA_HOSTNAME** environment variable, you must add it to the configuration file of the monitoring agent:
   - **On Windows:** Use the **Advanced > Edit Variables** option.
   - **On UNIX and Linux:** Add the variable to the `config/`*`product_code`*`.ini` and to `config/`*`product_code`*`.config` files.
   - **On z/OS:** Add the variable to the **RKANPAR** library, member K*product_code*ENV.
   - **On i5/OS:** Add the variable to the **QAUTOTMP/KMSPARM** library in member **KBBENV**.
3. Some monitoring agents (for example, the monitoring agent for MQ Series) do not reference the **CTIRA_HOSTNAME** environment variable to generate component names. Check the documentation for the monitoring agent that you are using for information on name generation. If necessary, contact IBM Software Support.

Agent troubleshooting

This section lists problems that might occur with agents.

This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Table 8. Agent problems and solutions

Problem	Solution
A configured and running instance of the monitoring agent is not displayed in the Tivoli Enterprise Portal, but other instances of the monitoring agent on the same system do appear in the portal.	Tivoli Monitoring products use Remote Procedure Call (RPC) to define and control product behavior. RPC is the mechanism that allows a client process to make a subroutine call (such as GetTimeOfDay or ShutdownServer) to a server process somewhere in the network. Tivoli processes can be configured to use TCP/UDP, TCP/IP, SNA, and SSL as the desired protocol (or delivery mechanism) for RPCs. "IP.PIPE" is the name given to Tivoli TCP/IP protocol for RPCs. The RPCs are socket-based operations that use TCP/IP ports to form socket addresses. IP.PIPE implements virtual sockets and multiplexes all virtual socket traffic across a single physical TCP/IP port (visible from the netstat command). A Tivoli process derives the physical port for IP.PIPE communications based on the configured, well-known port for the HUB Tivoli Enterprise Monitoring Server. (This well-known port or BASE_PORT is configured using the 'PORT:' keyword on the KDC_FAMILIES / KDE_TRANSPORT environment variable and defaults to '1918'.) The physical port allocation method is defined as (BASE_PORT + 4096*N) where N=0 for a Tivoli Enterprise Monitoring Server process and N={1, 2, ..., 15} for a non-Tivoli Enterprise Monitoring Server. Two architectural limits result as a consequence of the physical port allocation method: • No more than one Tivoli Enterprise Monitoring Server reporting to a specific Tivoli Enterprise Monitoring Server HUB can be active on a system image. • No more that 15 IP.PIPE processes can be active on a single system image. A single system image can support any number of Tivoli Enterprise Monitoring Server processes (address spaces) provided that each Tivoli Enterprise Monitoring Server on that image reports to a different HUB. By definition, there is one Tivoli Enterprise Monitoring Server HUB per monitoring Enterprise, so this architecture limit has been simplified to one Tivoli Enterprise Monitoring Server per system image. No more that 15 IP.PIPE processes or address spaces can be active on a single system image. With the first limit expressed above, this second limitation refers specifically to Tivoli Enterprise Monitoring Agent processes: no more that 15 agents per system image. This limitation can be circumvented (at current maintenance levels, IBM Tivoli Monitoring V6.1 Fix Pack 4 and later) if the Tivoli Enterprise Monitoring Agent process is configured to use EPHEMERAL IP.PIPE. (This is IP.PIPE configured with the 'EPHEMERAL:Y' keyword in the KDC_FAMILIES / KDE_TRANSPORT environment variable). There is no limitation to the number of ephemeral IP.PIPE connections per system image. If ephemeral endpoints are used, the Warehouse Proxy Agent is accessible from the Tivoli Enterprise Monitoring Server associated with the agents using ephemeral connections either by running the Warehouse Proxy Agent on the same computer or by using the Firewall Gateway feature. (The Firewall Gateway feature relays the Warehouse Proxy Agent connection from the Tivoli Enterprise Monitoring Server computer to the Warehouse Proxy Agent computer if the Warehouse Proxy Agent cannot coexist on the same computer.)
When you edit the configuration for an existing monitoring agent, the values displayed are not correct.	The original configuration settings might include non-ASCII characters. These values were stored incorrectly and result in the incorrect display. Enter new values using only ASCII characters.

Table 8. Agent problems and solutions (continued)

Problem	Solution
The Monitoring Agent for UNIX OS starts and displays in the Tivoli Enterprise Portal, but itmcmd/CandleAgent indicates that the agent has failed to start and is not running.	Check the `config/ux.ini` file for any blank lines. Delete them and restart the agent.
Attributes do not allow non-ASCII input in the situation editor.	None. Any attribute that does not include "(Unicode)" might support only ASCII characters. For example "Attribute (Unicode)" will support unicode but "Attribute" without "(Unicode)" might only support ASCII characters.
Changes made to the configuration of monitoring do not take effect.	Restart the monitoring agent so that your changes take effect.
(*For Monitoring Agent for UNIX OS agents running on AIX only*) When you use Logical Partitions (LPARs) on AIX 5.3 LPARs, the Monitoring Agent for UNIX OS fails to automatically detect some dynamic changes.	Restart the Monitoring Agent for UNIX OS so that it can detect the changes. For example, if the allocated memory for an LPAR changed, restart the agent.
You have installed the product manually, using an approach other than the one documented in the *IBM Tivoli Monitoring Installation and Setup Guide*. You need to confirm whether you have run the **SetPerm** command.	When you use the installation approach that is documented in the *IBM Tivoli Monitoring Installation and Setup Guide*, the **SetPerm** command is run automatically to set required permissions for the processes that IBM Tivoli Monitoring runs. When you do not use this approach the executables for the monitoring agent do not have the required privileges. Run the **SetPerm** command (which is located under *install_dir*/bin/ directory). The following example shows lists of agent binaries before and after they have the required privileges.

The monitoring agent support has the existing permissions:
`-rwxrwx---  1 itmuser  itmusers  32243 Sep 09 13:30 ifstat`
`-rwxrwx---  1 itmuser  itmusers  41045 Sep 09 13:30 kux_vmstat`
`-rwxrwx---  1 itmuser  itmusers  507562 Sep 09 13:30 kuxagent`
`-rwxrwx---  1 itmuser  itmusers  5772 Sep 09 13:30 kuxdstat`
`-rwxrwx---  1 itmuser  itmusers  42514 Sep 09 13:30 nfs_stat`
`-rwxr-sr-x  1 itmuser  system  128211 Sep 09 13:30 stat_daemon`
 	
The permissions for the Solaris2 monitoring agent are as follows:
`UID  r-s  r-x  r-x kuxagent  uid(0)  gid(3)`
 	
The user has the following permissions:
`real user  id(0)  effective user id(0)`
`real group  id(1)  effective group id(1)`
 	
If you have not run the **SetPerm** command, the following permissions are set:
`rwx rwx r-x uid(35008) gid(1111) kuxagent`
`rwx rwx r-x uid(35008) gid(1111) stat_daemon`
`rwx rwx r-x uid(35008) gid(1111) ifstat`
`rwx rwx r-x uid(35008) gid(1111) nfs_stat`
`rwx rwx r-x uid(35008) gid(1111) kuxagent`
 	
Note: If the log file has SUID, that means that you have run the **SetPerm** command.	

Table 8. Agent problems and solutions (continued)

Problem	Solution

Solaris agent is terminating unexpectedly.	Obtain the agent log file and verify whether it contains any of the following information in the log file:
	`bad_scan in server rpc`
	`bad_scan could be caused by nfsstate command output mismatch or NFS not active on this system***** nfs_stat terminating ****`
	`read 0 expected 248`
	`nsf-sd *** data collection terminated ***`
	
	If the log file has this type of information, see "Support information" on page 222.

The agent is installed and running normally. After rebooting the computer, where Tivoli Enterprise Monitoring Server was running, the agent is not online.	This problem can occur when the `root` user account is used to install and start the agent. Verify whether you have used the `root` user account to install the agent.To change the user account from `root` to some other user account, see Enabling the Monitoring Agent for UNIX OS to run as a nonroot user.
You want to have multiple instances of the same Monitoring Agent for UNIX OS running on the same system but talking to different Tivoli Enterprise Monitoring Server.	If you plan to install and run the Monitoring Agent for UNIX OS and Monitoring Agent for Linux OS agent on one computer, they can use the same network interface because they run as different processes. However, if you want to have two UNIX or two Linux agents on the same computer or want to run two instances of each agent, install two-network adapters. Each instance is configured for the host-specific address so they can be recognized in the configuration settings.
The Monitoring Agent for UNIX OS fails and the log file has the following message: `KUXDSTAT: Contact Customer Support disk performance table exceeded.`	This message is not related to the failure, so you can ignore it. If the agent is failing, search for a different cause. Further analyze the log to know whether the agent has terminated.
When you restart the system that hosts the Tivoli Enterprise Monitoring Server, the Monitoring Agent for UNIX OS does not start automatically. However, when you use CandleAgent start, the agent is starts and continues running.	If the agent does not connect to the Tivoli Enterprise Monitoring Server automatically, it means that you used the `root` user account to install and start the Monitoring Agent for UNIX OS. Most of the time, using the `root` account does not cause a problem, but the result is unpredictable. Check the **IBM Tivoli Monitoring `root` account to install and start the agent. To change the user account from `root` to another user account, see Enabling the Monitoring Agent for UNIX OS to run as a nonroot user.**

Table 8. Agent problems and solutions (continued)

Problem	Solution
The Monitoring Agent for UNIX OS (specifically the kuxagent process) uses a large amount of system resources.	In most cases, the problem occurs during the backup. Any one of the following scenarios can cause this problem. **The agent is running during the backup** After backing up, the agent is started during system startup. **Multiple agents are running at the same time.** The computer that hosts the Tivoli Enterprise Monitoring Server was rebooted and the agent has been installed by the `root` user account. **The agent is running during the backup** During the backup, some of the service might be interrupted or not be available or locked for some amount of time. While the backup process is going on, the Monitoring Agent for UNIX OS, which is running parallel, might wait for resources to be freed by the backup process. When the backup is completed and you are viewing the agent, high CPU at this point is expected, because the agent is in an uncertain state (backup usually stops several kernel services that could cause this state). For this reason, it is advisable to stop all agents before the backup run, because there might be lost information, file, or API connections. Stop the agent before the backup process starts. **The agent is started during system boot up:** If you use scripts to stop and start the agent, do not start the agent from an **init** process script when you restart the system. The computer that hosts the Tivoli Enterprise Monitoring Server was rebooted and the agent has been installed by the `root` user account. Verify whether the Monitoring Agent for UNIX OS log file has the following information: `Unable to find running Tivoli Enterprise Monitoring Server on CMSLIST`
You have two monitoring agents with the same name due to truncation.	Each name must be between 2 and 32 characters in length. Each agent name must be unique on the Tivoli Enterprise Monitoring Server. If the host name plus domain name length is greater than 32 characters multiple agents will conflict resulting in odd behavior (such as the agents appearing and disappearing from the Tivoli Enterprise Portal). Setting a unique host name resolves this issue. See "Unique names for monitoring components" on page 209 for more information.
Unicode filenames not properly being displayed in the File Information Viewer.	Due to incompatibilities in reading information from different language code pages, any file that has non-ascii text will not be properly displayed in the File Information viewer.
Unicode process names not properly being displayed in Process workspace.	Due to incompatibilities in reading information from different language code pages, any process that has non-ascii text will not be properly displayed in the Process Workspace.
On AIX systems, the Process.Base_Command attribute returns different data than expected.	The data should represent the name of the actual program being run. This is a current limitation on this platform.
If starting the monitoring agent on a Solaris system, you find you cannot access the file system, and then you have the same problem with the OS agent.	Add the non-administrator user used by the agent in the group having access to that file system.

Table 8. Agent problems and solutions (continued)

Problem	Solution
Query produces no historical data	If you use wildcards within a query, the value does not act as a wildcard against historical data. It acts as a value to compare against the values in the historical table, but it does act as a wildcard against the realtime data. For example, if you use .*(LongDirName/sleep).* in the historical collection configuration and use .*(LongDirName/sleep).* in a query as well, you will see real time data and historical data. But, if you use .*(LongDirName/sleep).* in the historical collection configuration, and then use .*(LongDirName/sle).* in the query, then you will see real time data only and no historical data.
Data collection for the AIX Defined Users attribute group degrades the performance of the Monitoring Agent for UNIX OS.	A variety of metrics have been ported from the AIX Premium agent to the Monitoring Agent for UNIX OS. New attribute groups include AIX AMS, AIX Defined Users, AIX Devices, AIX LPAR, AIX WPAR CPU, AIX WPAR File System, AIX WPAR Information, AIX WPAR Network, and AIX WPAR Physical Memory. Note that the data collection for the AIX Defined Users attribute group is by default disabled for performance reasons. You must set KUX_DEFINED_USERS=True in the `ux.ini` file to enable it. To view the variety of metrics from the AIX Premium agent, see Appendix B, “AIX Premium agent attributes,” on page 235. The minimum AIX requirements to collect the new metrics are • AIX53S = AIX 5.3 TL10 • AIX61F = AIX 6.1 TL3
Data collection of the metrics available from the kpx data provider degrades the performance of the Monitoring Agent for UNIX OS (or these metrics are not relevant in your environment).	By default, when running on AIX 5.3 or later, the kuxagent main spawns a new process. This process, aixdp_daemon, gathers all the metrics available from the kpx data provider and passes them back to kuxagent. By setting the environment variable KUX_AIXDP=false in the `ux.ini` file, an administrator has the option to prevent the aixdp_daemon process from starting. Specify this option when new metrics are not relevant or if performance issues arise. The variable's default value is true (that is, data collection is enabled by default). You can use the ITM V623 remote agent configuration feature to change the value.
Disk data collected by the UNIX OS agent does not match the data collected from the df (disk free) command output.	The df command provides the current total disk usage for all file systems accessible by the workstation. In contrast, disk data collected by the UNIX OS agent for the Space Used attributes includes the contribution from the reserved space, if any. As a result, these attributes might be reported as higher values than the df command's “Used" values which do not account for reserved space. In addition, disk data collected by the UNIX OS agent, expressed in percentages, is rounded up to the nearest integer. Percentages from the df command might not be rounded up to the nearest integer. Note that the ITM Size attribute does match the df command's “Total” output and equals the sum of Space Used and Space Available attributes.

Tivoli Enterprise Portal troubleshooting

Table 9 on page 216 lists problems that might occur with the Tivoli Enterprise Portal. This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Table 9. Tivoli Enterprise Portal problems and solutions

Problem	Solution
Historical data collection is unavailable because of incorrect queries in the Tivoli Enterprise Portal.	The column, Sort By, Group By, and First/Last functions are not compatible with the historical data collection feature. Use of these advanced functions will make a query ineligible for historical data collection. Even if data collection has been started, you cannot use the time span feature if the query for the chart or table includes any column functions or advanced query options (Sort By, Group By, First / Last). To ensure support of historical data collection, do not use the Sort By, Group By, or First/Last functions in your queries. See the *IBM Tivoli Monitoring Administrator's Guide* the Tivoli Enterprise Portal online Help for information on the Historical Data Collection function.
When you use a long process name in the situation, the process name is truncated.	Truncation of process names in the portal display is the expected behavior. 64 bytes is the maximum name length.

Troubleshooting for remote deployment

Table 10 lists problems that might occur with remote deployment. This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

This section describes problems and solutions for remote deployment and removal of agent software Agent Remote Deploy:

Table 10. Remote deployment problems and solutions

Problem	Solution
The removal of a monitoring agent fails when you use the remote removal process in the Tivoli Enterprise Portal desktop or browser.	This problem might happen when you attempt the remote removal process immediately after you have restarted the Tivoli Enterprise Monitoring Server. You must allow time for the monitoring agent to refresh its connection with the Tivoli Enterprise Monitoring Server before you begin the remote removal process.
Remote deploy fails.	Ensure that the OS agent is running as root.
Problems deploying agents through remote deployment.	If the user cannot deploy an agent without receiving an error, the deployed agent might require root permissions. Ensure that the deploying agent (UNIX agent binary - kuxagent) has root permissions by either starting it as root or assigning SETUID root permissions to the deploying agent.

Workspace troubleshooting

Table 11 shows problems that might occur with workspaces. This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Table 11. Workspace problems and solutions

Problem	Solution
You see the following message: `KFWITM083W Default link is disabled for the selected object; please verify link and link anchor definitions.`	You see this message because some links do not have default workspaces. Right-click the link to access a list of workspaces to select.

Table 11. Workspace problems and solutions (continued)

Problem	Solution
The name of the attribute does not display in a bar chart or graph view.	When a chart or graph view that includes the attribute is scaled to a small size, a blank space is displayed instead of a truncated name. To see the name of the attribute, expand the view of the chart until there is sufficient space to display all characters of the attribute's name.
A link to 'Process Resource' leads to a superseded workspace.	Use the provided workspace as a template for creating a custom workspace using the "Resource (621)" query. Then you can directly access the new workspace and see all the available processes on the system. You can create links to the workspace and also create links from the workspace pointing to other workspaces.
At the bottom of each view, you see the following Historical workspace KFWITM220E error: **Request failed during execution**.	Ensure that you configure all groups that supply data to the view. In the Historical Configuration view, ensure that data collection is started for all groups that supply data to the view.
You start collection of historical data but the data cannot be seen.	Managing options for historical data collection: • Basic historical data collection populates the Warehouse with raw data. This type of data collection is turned off by default. See Chapter 2, "Requirements for the monitoring agent," on page 7 for information on managing this feature including how to set the interval at which data is collected. By setting a more frequent interval for data collection you reduce the load on the system incurred every time data is uploaded. • You use the Summarization and Pruning monitoring agent to collect specific amounts and types of historical data. Be aware that historical data is not displayed until the Summarization and Pruning monitoring agent begins collecting the data. By default, this agent begins collection at 2 AM daily. At that point, data is visible in the workspace view. See the IBM Tivoli Monitoring Administrator's Guide to learn how to modify the default collection settings.
Inconsistent memory data shown in workspaces	For the Memory attribute group (displayed in the System Information workspace), the available real memory is obtained from the sysconf() system call and the swap space from the swapctl() system call, while in the Solaris Zones attribute group (displayed in the Solaris Zones workspace) both real memory and swap space usage come from the prstat command output. Differences are possible since prstat is not always reliable.
The CPU Share Pct definition shown in the workspace is inexact	The CPU shares are not equivalent to the percentage of CPU usage unless CPU demand is equal or greater to the available resources.
The meaning of Virtual Storage in the System Information workspace is not clear	By definition, Virtual Storage is obtained by summing up the Real Memory and the Paging Space, not considering the part of Paging Space that comes from reserved Real Memory for paging. In other words Virtual Storage = Real Memory + Paging Space - Real Memory for Paging.

Situation troubleshooting

This section provides information about both general situation problems and problems with the configuration of situations. See the *IBM Tivoli Monitoring Troubleshooting Guide* for more information about troubleshooting for situations.

Specific situation problems

Table 12 lists problems that might occur with specific situations.

Table 12. Specific situation problems and solutions

Problem	Solution
You want to change the appearance of situations when they are displayed in a Workspace view.	1. Right-click an item in the Navigation tree. 2. Select **Situations** in the pop-up menu. The Situation Editor window is displayed. 3. Select the situation that you want to modify. 4. Use the **Status** pull-down menu in the lower right of the window to set the status and appearance of the Situation when it triggers. **Note:** This status setting is not related to severity settings in IBM Tivoli Enterprise Console.
Situations are triggered in the Tivoli Enterprise Monitoring Server, but events for the situation are not sent to the Tivoli Enterprise Console server. The Tivoli Enterprise Monitoring Server is properly configured for event forwarding, and events for many other situations are sent to the event server.	This condition can occur when a situation is only monitoring the status of other situations. The event forwarding function requires an attribute group reference in the situation in order to determine the correct event class to use in the event. When the situation only monitors other situations, no attribute groups are defined and the event class cannot be determined. Because the event class cannot be determined, no event is sent. This is a limitation of the Tivoli Enterprise Monitoring Server event forwarding function. Situations that only monitor other situations do not send events to the event server.
Monitoring activity requires too much disk space.	Check the RAS trace logging settings that are described in "Setting RAS trace parameters" on page 204. For example, trace logs grow rapidly when you apply the **ALL** logging option.
A formula that uses mathematical operators is displayed to be incorrect. For example, if you were monitoring Linux, a formula that calculates when **Free Memory** falls under 10 percent of **Total Memory** does not work: `LT #'Linux_VM_Stats.Total_Memory' / 10`	This formula is incorrect because situation predicates support only logical operators. Your formulas cannot have mathematical operators. **Note:** The Situation Editor provides alternatives to math operators. Regarding the example, you can select **% Memory Free** attribute and avoid the need for math operators.
If you are running a Version 350 Monitoring Agent for UNIX OS and you choose to alter the views to include a Version 610 UNICODE attribute, be aware that data for this attribute is not displayed and you see a blank column in this view.	To enable Unicode and other features, upgrade the monitoring agent to IBM Tivoli Monitoring, Version 6.1.0.
The **Size** attribute in the **File Information** group of Monitoring Agent for UNIX OS provides file size metrics in bytes, and the resulting integers are so long that they are difficult to read.	Use the option to log size metrics in megabytes (MB). Future releases of this monitoring agent can provide the option to capture this metric in other units, such as KB.
You see the 'Unable to get attribute name' error in the Tivoli Enterprise Monitoring Server log after creating a situation.	Install the agent's application support files on the Tivoli Enterprise Monitoring Server, using the following steps: 1. Open the Manage Tivoli Enterprise Monitoring Services window. 2. Right-click the name of the monitoring server. 3. Select **Advanced > Add TEMS Application Support** in the pop-up menu. Add application support if any for any agent that is missing from the list. See the *IBM Tivoli Monitoring Installation and Setup Guide* for more information on adding application support.

Table 12. Specific situation problems and solutions (continued)

Problem	Solution
Events received at the Tivoli Enterprise Console server from IBM Tivoli Monitoring do not have values for all event attributes (slots) even though the values are visible in workspace views.	The problem is due to a limitation in the IBM Tivoli Monitoring interface code that generates Tivoli Enterprise Console events from situations. The situation results are provided in a chain of buffers of 3000 bytes each. The interface code currently extracts event information from only the first buffer. When situations or agent table data expands into a second buffer, this additional data is not examined, and it is not included in events sent to the Tivoli Enterprise Console server.
Tivoli Enterprise Console events from IBM Tivoli Monitoring 6.2 for IBM Tivoli Monitoring 5.x migrated situations receive parsing errors in the Tivoli Enterprise Console server.	Complete the following two steps: 1. Ensure that you have the IBM Tivoli Monitoring 6.2 Event Sync installed on your Tivoli Enterprise Console server. 2. Obtain updated baroc files from IBM Tivoli Monitoring 6.2 for the monitoring agent's events. Updated baroc files are on the Tivoli Enterprise Monitoring Server in the *CANDLEHOME*/CMS/TECLIB/itm5migr directory.
You are receiving Tivoli Business Systems Management events that cannot be associated due to application_oid and application_class not being set.	The problem is due to IBM Tivoli Monitoring 6.2 sending Tivoli Enterprise Console events for IBM Tivoli Monitoring 5.x migrated situations. These events are not able to set the cited slot values. Replace the *agent_name*_forward_tbsm_event_cb.sh script on the Tivoli Enterprise Console server with the version of this file from the Tivoli Enterprise Monitoring Server in the *CANDLEHOME*/CMS/TECLIB/itm5migr directory.

Situations you created using the File Pattern attribute group always raise alerts, sometimes unexpectedly.	A situation created using the File Pattern attribute group is always TRUE unless you incorporate the Match Count attribute into the formula. The Match Count attribute indicates the number of matches for the specified pattern in the specified file.
	
	For example, this formula always raises an alert:
	`IF VALUE Unix_File_Pattern.File_Name EQ '/path/filename'`
	`AND VALUE Unix_File_Pattern.Match_Pattern EQ 'pattern'`
	
	To remedy the of unexpected alerts raised by this type of situation, redefine the example by incorporating the Match Count attribute.
	`IF VALUE Unx_File_Pattern.File_Name EQ '/path/filename'`
	`AND VALUE Unix_File_Pattern.Match_Pattern EQ 'pattern' AND`
	`VALUE Unix_File_Pattern.Match_Count GT 0`

Problems with configuration of situations

Table 13 lists problems that might occur with situations.

This section provides information for troubleshooting for agents. Be sure to consult the *IBM Tivoli Monitoring Troubleshooting Guide* for more general troubleshooting information.

Table 13. Problems with configuring situations that you solve in the Situation Editor

Problem	Solution
Note: To get started with the solutions in this section, perform these steps: 1. Launch the Tivoli Enterprise Portal. 2. Click **Edit > Situation Editor**. 3. In the tree view, choose the agent whose situation you want to modify. 4. Choose the situation in the list. The Situation Editor view is displayed.	
The situation for a specific agent is not visible in the Tivoli Enterprise Portal.	Open the Situation Editor. Access the All managed servers view. If the situation is absent, confirm that application support for Monitoring Agent for UNIX OS has been added to the monitoring server. If not, add application support to the server, as described in the *IBM Tivoli Monitoring Installation and Setup Guide.*

Table 13. Problems with configuring situations that you solve in the Situation Editor (continued)

Problem	Solution
The monitoring interval is too long.	Access the Situation Editor view for the situation that you want to modify. Check the **Sampling interval** area in the **Formula** tab. Adjust the time interval as needed.
The situation did not activate at startup.	Manually recycle the situation as follows: 1. Right-click the situation and choose **Stop Situation**. 2. Right-click the situation and choose **Start Situation**. **Note:** You can permanently avoid this problem by placing a check mark in the **Run at Startup** option of the Situation Editor view for a specific situation.
The situation is not displayed.	Click the **Action** tab and check whether the situation has an automated corrective action. This action can occur directly or through a policy. The situation might be resolving so quickly that you do not see the event or the update in the graphical user interface.
An Alert event has not occurred even though the predicate has been properly specified.	Check the logs, reports, and workspaces.
A situation fires on an unexpected managed object.	Confirm that you have distributed and started the situation on the correct managed system.
The product did not distribute the situation to a managed system.	Click the **Distribution** tab and check the distribution settings for the situation.
The situation does not fire. Incorrect predicates are present in the formula that defines the situation. For example, the managed object shows a state that normally triggers a monitoring event, but the situation is not true because the wrong attribute is specified in the formula.	In the **Formula** tab, analyze predicates as follows: 1. Click the *fx* icon in the upper-right corner of the Formula area. The Show formula window is displayed. a. Confirm the following details in the **Formula** area at the top of the window: • The attributes that you intend to monitor are specified in the formula. • The situations that you intend to monitor are specified in the formula. • The logical operators in the formula match your monitoring goal. • The numerical values in the formula match your monitoring goal. b. (*Optional*) Click the **Show detailed formula** check box in the lower left of the window to see the original names of attributes in the application or operating system that you are monitoring. c. Click **OK** to dismiss the Show formula window. 2. (*Optional*) In the Formula area of the **Formula** tab, temporarily assign numerical values that will immediately trigger a monitoring event. The triggering of the event confirms that other predicates in the formula are valid. **Note:** After you complete this test, you must restore the numerical values to valid levels so that you do not generate excessive monitoring data based on your temporary settings.

Table 14. Problems with configuration of situations that you solve in the Workspace area

Problem	Solution
Situation events are not displayed in the Events Console view of the workspace.	Associate the situation with a workspace. **Note:** The situation does not need to be displayed in the workspace. It is sufficient that the situation be associated with any workspace.

Table 14. Problems with configuration of situations that you solve in the Workspace area (continued)

Problem	Solution
You do not have access to a situation.	**Note:** You must have administrator privileges to perform these steps. 1. Select **Edit** > **Administer Users** to access the Administer Users window. 2. In the Users area, select the user whose privileges you want to modify. 3. In the Permissions tab, Applications tab, and Navigator Views tab, select the permissions or privileges that correspond to the user's role. 4. Click **OK**.
A managed system seems to be offline.	1. Select Physical View and highlight the Enterprise Level of the navigator tree. 2. Select **View** > **Workspace** > **Managed System Status** to see a list of managed systems and their status. 3. If a system is offline, check network connectivity and status of the specific system or application.

Table 15. Problems with configuration of situations that you solve in the Manage Tivoli Enterprise Monitoring Services window

Problem	Solution
After an attempt to restart the agents in the Tivoli Enterprise Portal, the agents are still not running.	For UNIX, NetWare, or Windows, log on to the applicable system and perform the appropriate queries.
The Tivoli Enterprise Monitoring Server is not running.	Check the system status and check the appropriate IBM Tivoli Monitoring logs.
The managed objects you created are firing on incorrect managed systems.	Check the managed system distribution on both the situation and the managed object settings sheets.

Take Action command troubleshooting

Table 16 lists general problems that might occur with Take Action commands. When each Take Action command runs it generates the log file listed in Table 5 on page 202. This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Table 16. Take Action commands problems and solutions

Problem	Solution
Take Action commands might require several minutes to complete.	Allow several minutes. If you do not see a pop-up message advising you of completion, try to run the command manually. If you are unable to perform the Take Action command manually, see *IBM Tivoli Monitoring Troubleshooting Guide* for general information on troubleshooting the Take Action command.

Troubleshooting for UNIX

Table 17 on page 222 lists problems that might occur on the system or application that you are monitoring. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Table 17. Paging and memory issues for a system administrator to consider

Problem	Solution
Paging space allotted needs to be increased.	A system needs to have about two times its total memory size. However, paging space depends on the programs that are running. If the site has many small programs that run to completion quickly, then only one times the total memory size might be required. If the site runs large programs that run for hours or days at a time, then more paging space is required.
Paging space allocation needs to be more accurate on a UNIX system.	You can test the allocation of paging space by creating a situation that monitors Active Virtual Memory. Active Virtual Memory closely matches how much paging space is being used. When the system uses all the paging space, the operating system terminates processes that ask for more. To create a situation that monitors active virtual memory: 1. Use the UNIX detail view to obtain the Total Virtual Memory, and to compute 90% and 95% of the Total Virtual Memory. 2. When the Active Virtual Memory is equal to 90%, of the Total Virtual Memory this is a Yellow light condition. When the Active Virtual Memory is equal to 95% of the Total Virtual Memory this is a Red light condition. In response to this test, the local system administrator can increase the percentages or lower them, as appropriate.
The Free Memory value seems too small.	The System Report, Free Memory column displays how much free memory is available at the current time. This number is normally small. However, you must take action if this number is zero and remains zero for a long period of time. On AIX systems a small number means that the operating system is doing an efficient job at managing the memory of the system. If this number is very large, the system is not busy and has more RAM than required.

Support information

If you have a problem with your IBM software, you want to resolve it quickly. IBM provides the following ways for you to obtain the support you need:

Online

Go to the IBM Software Support site at http://www.ibm.com/software/support/probsub.html and follow the instructions.

IBM Support Assistant

The IBM Support Assistant (ISA) is a free local software serviceability workbench that helps you resolve questions and problems with IBM software products. The ISA provides quick access to support-related information and serviceability tools for problem determination. To install the ISA software, go to http://www.ibm.com/software/support/isa.

Accessing terminology online

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/software/globalization/terminology

Accessing publications online

IBM posts publications for this and all other Tivoli products, as they become available and whenever they are updated, to the Tivoli Documentation Central Web site at http://www.ibm.com/tivoli/documentation.

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 http://www.elink.ibmlink.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.ibmlink.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.

Tivoli technical training

For Tivoli technical training information, refer to the following IBM Tivoli Education Web site at http://www.ibm.com/software/tivoli/education.

Tivoli user groups

Tivoli user groups are independent, user-run membership organizations that provide Tivoli users with information to assist them in the implementation of Tivoli Software solutions. Through these groups, members can share information and learn from the knowledge and experience of other Tivoli users. Tivoli user groups include the following members and groups:

- 23,000+ members
- 144+ groups

Access the link for the Tivoli Users Group at https://community.ibm.com/community/user/imwuc/home.

Appendix A. Upgrading for warehouse summarization

The Monitoring Agent for UNIX OS made changes to the warehouse collection and summarization characteristics for some agent attribute groups. These changes correct and improve the way warehouse data is summarized, producing more meaningful historical reports. This appendix explains those changes and the implications to your warehouse collection and reporting.

Note: This upgrade is not available for upgrading from IBM Tivoli Monitoring V6.2.1 to V6.2.3. Use these scripts only when you upgrade to V6.2.1 (or later) from a prior version, or to V6.2.3 Fix Pack 1 (or later) from a prior version of 6.2.3.

Warehouse summarization is controlled on a per-table basis. How the rows in each table are summarized is determined by a set of attributes in each table that are designated as primary keys. There is always one primary key representing the monitored resource, and data is minimally summarized based on this value. For all agents, this primary key is represented internally by the column name, ORIGINNODE; however, the external attribute name varies with each monitoring agent.

One or more additional primary keys are provided for each attribute group to further refine the level of summarization for that attribute group. For example, in an OS agent disk attribute group, a primary key might be specified for the logical disk name that allows historical information to be reported for each logical disk in a computer.

Tables in the warehouse

For a monitoring agent, there are two main types of warehouse tables:

- Raw tables:

  These tables contain the raw information reported by a monitoring agent and written to the warehouse by the Warehouse Proxy agent. Raw tables are named for the attribute group that they represent, for example, N_F_S_and_R_P_C_Statistics.

- Summary tables:

  These tables contain summarized information based on the raw tables and written to the warehouse by the Summarization and Pruning agent. Summarization provides aggregation results over various reporting intervals, for example, hours, days, and so on. Summary table names are based on the raw table name with an appended suffix, for example, N_F_S_and_R_P_C_Statistics_H, N_F_S_and_R_P_C_Statistics_D, and so on.

Effects on summarized attributes

When tables are summarized in the warehouse, the summary tables and summary views are created to include additional columns to report summarization information. Table 18 on page 226 contains a list of the time periods and the suffixes for the summary tables and views.

Table 18. Time periods and suffixes for summary tables and views

Data collection time period	Summary table suffixes	Summary view suffixes
Hourly	_H	_HV
Daily	_D	_DV
Weekly	_W	_WV
Monthly	_M	_MV
Quarterly	_Q	_QV
Yearly	_Y	_YV

Table 19 shows the expansion to summary columns of some of the most commonly used attribute types.

Table 19. Additional columns to report summarization information

Attribute name	Aggregation type	Additional summarization columns
MyGauge	GAUGE	MIN_MyGauge MAX_MyGauge SUM_MyGauge AVG_MyGauge
MyCounter	COUNTER	TOT_MyCounter HI_MyCounter LO_MyCounter LAT_MyCounter
MyProperty	PROPERTY	LAT_Property

These additional columns are provided only for attributes that are not primary keys. In the cases when an existing attribute is changed to be a primary key, the Summarization and Pruning agent no longer creates summarization values for the attributes, but the previously created column names remain in the table with any values already provided for those columns. These columns cannot be deleted from the warehouse database, but as new data is collected, these columns will not contain values. Similarly, when the primary key for an existing attribute has its designation removed, that attribute has new summarization columns automatically added. As new data is collected, it is used to populate these new column values, but any existing summarization records do not have values for these new columns.

The overall effect of these primary key changes is that summarization information is changing. If these changes result in the old summarization records no longer making sense, you can delete them. As a part of warehouse upgrade, summary views are dropped. The views will be recreated by the Summarization and Pruning agent the next time it runs. Dropping and recreating the views ensure that they reflect the current table structure.

Upgrading your warehouse with limited user permissions

The IBM Tivoli Monitoring warehouse agents (Warehouse Proxy and Summarization and Pruning agents) can dynamically adjust warehouse table definitions based on attribute group and attribute information being loaded into the warehouse. These types of table changes must be done for this monitoring agent for one or both of the following conditions:

- The monitoring agent has added new attributes to an existing attribute group and that attribute group is included in the warehouse.
- The monitoring agent has added a new attribute group and that attribute group is included in the warehouse.

For the warehouse agents to automatically modify the warehouse table definitions, they must have permission to alter warehouse tables. You might not have granted these agents these permissions, choosing instead to manually define the raw tables and summary tables needed for the monitoring agents. Or, you might have granted these permissions initially, and then revoked them after the tables were created.

You have two options to effect the required warehouse table changes during the upgrade process:

- Grant the warehouse agents temporary permission to alter tables

  If using this option, grant the permissions, start historical collection for all the desired tables, allow the Warehouse Proxy agent to add the new data to the raw tables, and allow the Summarization and Pruning agent to summarize data for all affected tables. Then, remove the permission to alter tables
- Make the warehouse table updates manually

  If using this option, you must determine the table structures for the raw and summary tables. If you manually created the tables in the earlier warehouse definition, you already have a methodology and tools to assist you in this effort. You can use a similar technique to update and add new tables for this warehouse migration.

  For a method of obtaining raw table schema, refer to the IBM Redbook,*Tivoli Management Services Warehouse and Reporting*, January 2007, SG24-7290. The chapter that explains warehouse tuning includes a section on creating data tables manually.

Types of table changes

The following types of table changes affect warehouse summarization:

Case 1 - New attribute added to an attribute group and defined as a primary key.

Case 2 - Existing attribute defined as a primary key or had primary key designation removed.

Case 3 - Moving some tables from 4K tablespaces to 8K tablespaces when using DB2 as the warehouse database.

Case 1 and Case 2 are primary key changes. In both cases, new summarization records will not match existing summarized data:

- A new attribute was added to an attribute group and that attribute was defined as a primary key:

  New summarization records will provide more accurate summarization or greater granularity than previous records. Existing summarization records are still available but contain less granular detail if default values are not assigned for the new primary keys.
- An existing attribute was defined as a primary key or the primary key designation was removed:

  If a new key was added, then the new summarization records will provide more accurate summarization or greater granularity than previous records. If a key was removed, then the new summarization records will provide less granularity

than previous records, but with the intent of providing more meaningful summarization. Existing summarization records are still available.

Case 3 requires that you move some tables from 4K tablespaces to 8K tablespaces when using DB2 as the warehouse database to avoid errors during summarization and pruning processing.

Table summary

Table 20 provides information to help you determine the effects of primary key and warehouse changes for this monitoring agent. The table shows each attribute group, the current primary keys (in addition to ORIGINNODE) for the attribute group, primary keys that were removed, and whether this table is being included in warehouse reporting.

Table 20. Primary key and warehouse changes for the Monitoring Agent for UNIX OS

Attribute group	Current primary keys	Removed primary keys	Warehoused
Disk_Performance	Disk_Name_U		Yes
Disk	Name_U		Yes
File_Information	File_U Path_U		No
Machine_Information			Yes
N_F_S_and_R_P_C_Statistics	NFS_Version		Yes
Network	Network_Interface_Name		Yes
Process	Process_ID		Yes
SMP_CPU	CPU_ID		Yes
Solaris_Zones	Path Zone_ID Name		Yes
System SP2System			Yes
UNIX_IP_Address	IP_Address Network_Interface_Name		Yes
UNIX_All_Users	User_ID Name		Yes
UNIX_File_Comparison			No
UNIX_File_Pattern			No
UNIX_Group	Group_ID Group_Name		Yes
UNIX_Memory			Yes
UNIX_Ping	Target_Host		Yes
UNIX_Print_Queue	Device_Name Print_Queue_Name		Yes
User	Login_Name_U		Yes

Upgrading your warehouse for primary key and tablespace changes

Upgrading your warehouse includes making the following types of changes:

- Case 1 - New attribute is added and is designated as a primary key
  - New attribute and a default value must be added to the raw table and the summarization tables.

    If the attribute group name is not too large for the underlying database, the table name corresponds to the attribute group name. If the attribute group name is too long, a short name is used. The mapping of attribute group names to table names is stored in the WAREHOUSEID table.

  - Case-1 scripts that perform the following actions are provided to assist in this change:
    - Alter existing raw tables
    - Alter existing summary tables
    - Drop existing summary views
  - These changes must be done before the monitoring agent is started and begins exporting data to the Warehouse Proxy agent.
- Case-2 - Existing attributes are changed to either add or remove primary key designation.
  - Existing data is of limited value and should be deleted.
  - Case-2_Truncate scripts that perform the following actions are provided to assist in this change:
    - Remove all records from existing summary tables, preserving existing table definitions
    - Delete the raw data marker allowing raw data to be resummarized
  - Case-2_Drop scripts that perform the following actions are provided to assist in this change:
    - Drop existing summary views
    - Drop existing summary tables
    - Delete the raw data marker allowing raw data to be resummarized
  - These changes are optional, but result in more accurate summarized information.
- Case 3 - Move tables from 4K tablespace to 8K tablespace for selected agents
  - Special processing for selected agents, to move tables from a 4K tablespace to an 8K tablespace.
  - Individual scripts are provided for each summary table to be changed.

Affected attribute groups and supporting scripts

Table 21 shows the attribute groups and summary tables affected for this monitoring agent, the names of the SQL scripts provided to assist in the upgrade process, the types of warehouse databases for which the scripts must be run, and the types of changes (cases) to which the scripts apply.

Table 21. Scripts for affected attribute groups and summary tables for the Monitoring Agent for UNIX OS

Attribute group or summary table	File	DB2	Oracle	MS SQL Server	Case 1	Case 2	Case 3
N_F_S_and_R_P_C _Statistics_D	kux_61migr_N_F_S_and_R_P_C _Statistics_D.sql	X			X		X
N_F_S_and_R_P_C _Statistics_H	kux_61migr_N_F_S_and_R_P_C _Statistics_H.sql	X			X		X
N_F_S_and_R_P_C _Statistics_M	kux_61migr_N_F_S_and_R_P_C _Statistics_M.sql	X			X		X
N_F_S_and_R_P_C _Statistics_Q	kux_61migr_N_F_S_and_R_P_C _Statistics_Q.sql	X			X		X
N_F_S_and_R_P_C _Statistics_W	kux_61migr_N_F_S_and_R_P_C _Statistics_W.sql	X			X		X
N_F_S_and_R_P_C _Statistics_Y	kux_61migr_N_F_S_and_R_P_C _Statistics_Y.sql	X			X		X

Table 21. Scripts for affected attribute groups and summary tables for the Monitoring Agent for UNIX OS (continued)

Attribute group or summary table	File	DB2	Oracle	MS SQL Server	Case 1	Case 2	Case 3
N_F_S_and_R_P_C _Statistics	kux_61migr_UNIX_OS_Agent_ Case-1.sql	X	X	X	X		
System_D	kux_61migr_System_D.sql	X					X
System_H	kux_61migr_System_H.sql	X					X
System_M	kux_61migr_System_M.sql	X					X
System_Q	kux_61migr_System_Q.sql	X					X
System_W	kux_61migr_System_W.sql	X					X
System_Y	kux_61migr_System_Y.sql	X					X
Note: The raw table System does not need to be modified even though the corresponding summary tables are modified.							

The following types of warehouse objects are affected by these scripts. Review the scripts before running them:

- Case-1.sql

  These scripts affect raw tables, summary tables, and summary views.
- Case-2_Drop.sql

  These scripts affect the summary tables, summary views, and the Summarization and Pruning agent WAREHOUSEMARKER table.
- Case-2_Truncate.sql

  These scripts affect the summary tables and the Summarization and Pruning agent WAREHOUSEMARKER table.

Procedures

The warehouse can be hosted on any of three databases: DB2, Oracle, or Microsoft SQL Server. There are different sets of script files for each type of database. These scripts are provided as part of the monitoring agent Tivoli Enterprise Portal Server support file installation. After installing the Tivoli Enterprise Portal Server support files for the monitoring agent, the files are located on the Tivoli Enterprise Portal Server computer in *install_dir*/CNPS/SQLLIB/WAREHOUSE. There is a subdirectory for each type of database: DB2 for DB2, Oracle for Oracle, and SQLServer for Microsoft SQL Server.

The scripts provide commands for all affected tables and views. If you do not have summarization enabled for some periods, for example, quarterly or yearly, you will not have the corresponding summary tables (_Q, _Y) and summary views (_QV, _YV) in your warehouse database. If you run the scripts that are provided, the database reports errors for these missing objects. The scripts continue to run the remaining commands. Similarly, if you rerun the scripts, all commands are attempted. If the objects do not exist, or the command cannot be run (especially for the ALTER commands), the scripts continue processing the remaining commands.

DB2 warehouse database procedure

1. Stop *all* running Warehouse Proxy agent instances and the Summarization and Pruning agent.
2. Back up your warehouse database.

3. Copy the scripts from the Tivoli Enterprise Portal Server in one of the following directories to a temporary directory on the system where the warehouse database is located:
   - Windows:

     install dir\CNPS\SQLLIB\WAREHOUSE\DB2
   - UNIX and Linux:

     install dir/*arch*/cq/sqllib/WAREHOUSE/DB2
4. On the system where the warehouse database is located, change to the directory where you placed the script files in Step 3. Then, connect to the warehouse database through the DB2 command line with a user ID that has the authorization to alter and load tables and drop views. Run commands based on the following example to connect, set the schema, and save the script to an output file:

   ```
   db2 connect to WAREHOUS user ITMUSER using ITMPASS
   db2 set current schema="ITMUSER"
   db2 -tv -z log/script.sql.log -f script.sql
   ```

   These parameters are used in the example:
   - WAREHOUS is the database name.
   - ITMUSER is the user name used by the Warehouse Proxy agent.
   - ITMPASS is the password used by the Warehouse Proxy agent.
   - *script.sql* is the name of the script file. See Table 21 on page 229 for the script file names.
   - *script.sql*.log is the name of the output file.

   Notes: You might receive error messages such the following from DB2:
   - `SQL0204N "schema name.table name" is an undefined name. SQLSTATE=42704`

     This message indicates that the table named *table name* does not exist and cannot be altered or dropped. This happens if you do not have warehousing or summarization enabled for the given table. For example if you only have hourly and daily summarization enabled, you see this message for the weekly, monthly, quarterly, and yearly summarization tables because these tables do not exist.
   - `SQL3304N The table does not exist.`

     This message indicates that the table does not exist and cannot be loaded. This happens if you do not have warehousing or summarization enabled for the given table. For example if you only have hourly and daily summarization enabled, you see this message for the weekly, monthly, quarterly, and yearly summarization tables because these tables do not exist.

Oracle warehouse database procedure

1. Stop *all* running Warehouse Proxy agent instances and the Summarization and Pruning agent.
2. Back up your warehouse database.
3. Copy the scripts from The Tivoli Enterprise Portal Server in one of the following directories to a temporary directory on the system where the warehouse database is located:
   - Windows

     install dir\CNPS\SQLLIB\WAREHOUSE\Oracle
   - UNIX and Linux

install dir/*arch*/cq/sqllib/WAREHOUSE/Oracle

4. On the system where the warehouse database is located, change to the directory where you placed the script files in Step 3. Then, connect to the warehouse database through the Oracle command line with the same user that the Warehouse Proxy agent uses to connect to the warehouse, and run the script. To run the script, the user ID must have authorization to alter tables and drop views, or to drop tables when using Case 2 Drop, or truncate tables when using Case 2 Truncate. The output is saved to a file named *script name*.log. Run the following command:

   ```
   sqlplus ITMUSER/ITMPASS@WAREHOUS @script.sql
   ```

   These parameters are used in the example:

   - WAREHOUS is the connect identifier.
   - ITMUSER is the user name used by the Warehouse Proxy agent.
   - ITMPASS is the password used by the Warehouse Proxy agent.
   - *script.sql* is the name of this script file. See Table 21 on page 229 for the script file names.

   Note: You might receive error messages such as the following from Oracle:

   ```
   ORA-00942: table or view does not exist
   ```

   This message indicates that the table does not exist and cannot be altered, dropped, or truncated. This happens if you do not have warehousing or summarization enabled for the given table. For example if you only have hourly and daily summarization enabled, you see this message for the weekly, monthly, quarterly, and yearly summarization tables because these tables do not exist.

MS SQL warehouse database procedure

1. Stop *all* running Warehouse Proxy agent instances and the Summarization and Pruning agent.
2. Back up your warehouse database.
3. Copy the scripts from the Tivoli Enterprise Portal Server in the one of the following directories to a temporary directory on the system where the warehouse database is located:
   - Windows:

     install dir\CNPS\SQLLIB\WAREHOUSE\SQLServer
   - UNIX and Linux:

     install dir/*arch*/cq/sqllib/WAREHOUSE/SQLServer
4. On the system where the warehouse database is located, change to the directory where you placed the script files in Step 3. Then, connect to the warehouse database through the SQL Server command line with the same user that the Warehouse Proxy agent uses to connect to the warehouse, and run the script. To run the script, the user ID must have authorization to alter tables and drop views, or to drop tables when using Case 2 Drop, or truncate tables when using Case 2 Truncate. The output is saved to a file named *script name*.log. Run the following command:

   ```
   osql -I -S SQLHOST[\SQLINST] -U ITMUSER -P ITMPASS -d WAREHOUS
      -m-1 -n -o log/script.sql.log -i script.sql
   ```

   These parameters are used in the example:

   - WAREHOUS is the database name.
   - ITMUSER is the user name used by the Warehouse Proxy agent.

- ITMPASS is the password used by the Warehouse Proxy agent.
- *script.sql* is the name of this script file.
- SQLHOST is the SQL server name.
- SQLINST is the optional SQL instance name.

Note: You might receive error messages from the SQL Server such as the following: `Msg 4902, Level 16, State 1, Server ENTERPRISE, Line 1 Cannot find the object "table name" because it does not exist or you do not have permissions.`

This message indicates that the table named *table name* does not exist and cannot be dropped or truncated. This happens if you do not have warehousing or summarization enabled for the given table. For example if you only have hourly and daily summarization enabled, you see this message for the weekly, monthly, quarterly, and yearly summarization tables because these tables do not exist.

Appendix B. AIX Premium agent attributes

The Monitoring Agent for UNX OS includes a variety of metrics from the AIX Premium agent. Each of the following tables maps the AIX Premium agent's attributes of a specific attribute group to their corresponding UNIX OS agent's attributes.

Table 22. AMS_Pool attribute group (table name: KPX_AMS_POOL)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
AMS_Mem_Loaned	AML	AIX_AMS	UNIXAMS	AMS_Mem_Loaned	AML
AMS_Memory_Ent_InUse	AMEI	AIX_AMS	UNIXAMS	AMS_Memory_Ent_InUse	AMEI
AMS_Memory_Entitlement	AME	AIX_AMS	UNIXAMS	AMS_Memory_Entitlement	AME
AMS_Mode	AMS _MODE	AIX_AMS	UNIXAMS	AMS_Mode	AMS _MODE
AMS_Physical_Mem	APM	AIX_AMS	UNIXAMS	AMS_Physical_Mem	APM
AMS_Pool_ID	API	AIX_AMS	UNIXAMS	AMS_Pool_ID	API
AMS_Pool_Size	APS	AIX_AMS	UNIXAMS	AMS_Pool_Size	APS
Hypervisor_Page_Ins	HPI	AIX_AMS	UNIXAMS	Hypervisor_Page_Ins	HPI
Hypervisor_Page_Ins_Time	HPIT	AIX_AMS	UNIXAMS	Hypervisor_Page_Ins_Time	HPIT
System_Name	NODE	AIX_AMS	UNIXAMS	System_Name	ORIGIN NODE
Timestamp	TIME STAMP	AIX_AMS	UNIXAMS	Timestamp	TIME STAMP

Table 23. CPU_Detail attribute group (table name: KPX_CPU_DETAIL)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Context_Switches_per_Sec	CSPS	SMP_CPU	UNIXCPU	Context_Switches_per_Sec	CSPS
Logical_Context_Switches	LCS	SMP_CPU	UNIXCPU	Logical_Context_Switches	LCS
Physical_Consumption	PC	SMP_CPU	UNIXCPU	Physical_Consumption	PC

Table 24. CPU_Summary attribute group (table name: KPX_CPU_SUMMARY)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Donated_Busy_Cycles_Pct	DBCP	AIX_LPAR	UNIXLPAR	Donated_Busy_Cycles_Pct	DBCP
Donated_Idle_Cycles_Pct	DICP	AIX_LPAR	UNIXLPAR	Donated_Idle_Cycles_Pct	DICP

Table 24. CPU_Summary attribute group (table name: KPX_CPU_SUMMARY) (continued)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Donation_Enablement	DE	AIX_LPAR	UNIXLPAR	Donation_Enablement	DE
Hypervisor_Calls	HC	AIX_LPAR	UNIXLPAR	Hypervisor_Calls	HC
Number_of_CPUs	NOC	System	UNIXOS	Number_of_CPUs	NOC
Physical_Consumption	PC	System	UNIXOS	Physical_Consumption	PC
Stolen_Busy_Cycles_Pct	SBCP	System	UNIXOS	Stolen_Busy_Cycles_Pct	SBCP
Stolen_Idle_Cycles_Pct	SICP	System	UNIXOS	Stolen_Idle_Cycles_Pct	SICP
System_Software_Version	SSV	System	UNIXOS	System_Software_Version	SSV
Time_Spent_in_Hypervisor _Pct	TSIHP	System	UNIXOS	Time_Spent_in_Hypervisor _Pct	TSIHP

Table 25. Defined_Users attribute group (table name: KPX_DEFINED USERS)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Account_Locked	AL	AIX _Defined _Users	UNIXDUSERS	Account_Locked	AL
Expires	EXPIRES	AIX _Defined _Users	UNIXDUSERS	Expires	EXPIRES
Loginretries	L	AIX _Defined _Users	UNIXDUSERS	Loginretries	L
Roles	ROLES	AIX _Defined _Users	UNIXDUSERS	Roles	ROLES
System_Name	NODE	AIX _Defined _Users	UNIXDUSERS	System_Name	ORIGIN NODE
Timestamp	TIME STAMP	AIX _Defined _Users	UNIXDUSERS	Timestamp	TIME STAMP
User_Name	USER _NAME	AIX _Defined _Users	UNIXDUSERS	User_Name	USER _NAME

Table 26. Devices attribute group (table name: KPX_DEVICES)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Class	CLASS	AIX _Devices	UNIXDEVIC	Dclass	DCLASS
Name	NAME	AIX _Devices	UNIXDEVIC	Name	NAME
Parent	PARENT	AIX _Devices	UNIXDEVIC	Parent	PARENT
State	STATE	AIX _Devices	UNIXDEVIC	State	STATE
System_Name	NODE	AIX _Devices	UNIXDEVIC	System_Name	ORIGIN NODE
Timestamp	TIME STAMP	AIX _Devices	UNIXDEVIC	Timestamp	TIME STAMP
Type	TYPE	AIX _Devices	UNIXDEVIC	Type	TYPE

Table 27. Disks attribute group (table name: KPX_DISKS)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Avg_ServiceQ_Size	ASS	Disk_ Performance	UNIXDPERF	Avg_ServiceQ_Size	ASS
Parent	PARENT	Disk_ Performance	UNIXDPERF	Parent	PARENT
ServiceQ_Full_per_Sec	SFPS	Disk_ Performance	UNIXDPERF	ServiceQ_Full_per_Sec	SFPS
Transfers_KB_per_Sec	TKPS	Disk_ Performance	UNIXDPERF	Transfers_KB_per_Sec	TKPS
Type	TYPE	Disk_ Performance	UNIXDPERF	Type	TYPE

Table 28. File_System attribute group (table name: KPX_FILE_SYSTEMS)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Volume_Group_Name	VGN	Disk	UNIXDISK	Volume_Group_Name	VGN

Table 29. Logical_Partition attribute group (table name: KPX_LOGICAL_PARTITION)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Available_CPU_Units _in_Pool	ACUIP	AIX_LPAR	UNIXLPAR	Available_CPU_Units _in_Pool	ACUIP
Available_CPUs_in_Pool	ACIP	AIX_LPAR	UNIXLPAR	Available_CPUs_in_Pool	ACIP
Busy_Pct	BUSY _PCT	AIX_LPAR	UNIXLPAR	Busy_Pct	BUSY _PCT
Capacity_Weight	CW	AIX_LPAR	UNIXLPAR	Capacity_Weight	CW
Capped_Mode	CM	AIX_LPAR	UNIXLPAR	Capped_Mode	CM
CPU_Entitlement	CE	AIX_LPAR	UNIXLPAR	CPU_Entitlement	CE
Entitlement	E	AIX_LPAR	UNIXLPAR	Entitlement	E
Entitlement_Pct	EP	AIX_LPAR	UNIXLPAR	Entitlement_Pct	EP
Entitlement_Used_Pct	EUP	AIX_LPAR	UNIXLPAR	Entitlement_Used_Pct	EUP
Hostname	HOST NAME	AIX_LPAR	UNIXLPAR	Hostname	HOST NAME
Last_Machine_ID	LMI	AIX_LPAR	UNIXLPAR	Last_Machine_ID	LMI
LPAR_Number	LN	AIX_LPAR	UNIXLPAR	LPAR_Number	LN
Machine_ID	MACHINE _ID	AIX_LPAR	UNIXLPAR	Machine_ID	MACHINE _ID
Max_CPU_Cap_Used_Pct	MCCUP	AIX_LPAR	UNIXLPAR	Max_CPU_Cap_Used_Pct	MCCUP
Maximum_Pool_Capacity	MPC0	AIX_LPAR	UNIXLPAR	Maximum_Pool_Capacity	MPC0
Number_of_Logical_CPUs	NOLC	AIX_LPAR	UNIXLPAR	Number_of_Logical_CPUs	NOLC
Number_of_Physical_CPUs	NOPC	AIX_LPAR	UNIXLPAR	Number_of_Physical_CPUs	NOPC
Number_of_Physical_CPUs _in_Shared_Pool	NOPCISP	AIX_LPAR	UNIXLPAR	Number_of_Physical_CPUs _in_Shared_Pool	NOPCISP
Number_of_Virtual_CPUs	NOVC	AIX_LPAR	UNIXLPAR	Number_of_Virtual_CPUs	NOVC
Phantom_Interrupts	PI	AIX_LPAR	UNIXLPAR	Phantom_Interrupts	PI
Phys_Busy_Pct	PBP	AIX_LPAR	UNIXLPAR	Phys_Busy_Pct	PBP
Physical_CPU_Size_of _Shared_Pool	PCSOSP	AIX_LPAR	UNIXLPAR	Physical_CPU_Size_of _Shared_Pool	PCSOSP
Physical_CPU_Units_Used	PCUU	AIX_LPAR	UNIXLPAR	Physical_CPU_Units_Used	PCUU
Pool_Entitlement	PE	AIX_LPAR	UNIXLPAR	Pool_Entitlement	PE
PoolID	POOLID	AIX_LPAR	UNIXLPAR	PoolID	POOLID
Shared_Mode	SM	AIX_LPAR	UNIXLPAR	Shared_Mode	SM
SMT_Mode	SMT _MODE	AIX_LPAR	UNIXLPAR	SMT_Mode	SMT _MODE
SMT_Threads	ST	AIX_LPAR	UNIXLPAR	SMT_Threads	ST
System_Name	NODE	AIX_LPAR	UNIXLPAR	System_Name	ORIGIN NODE
Timestamp	TIME STAMP	AIX_LPAR	UNIXLPAR	Timestamp	TIME STAMP
Total_Used_Pct	TUP	AIX_LPAR	UNIXLPAR	Total_Used_Pct	TUP

Table 29. Logical_Partition attribute group (table name: KPX_LOGICAL_PARTITION) (continued)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Unallocated_CPU_In_Pool	UCIP	AIX_LPAR	UNIXLPAR	Unallocated_CPU_In_Pool	UCIP
Uptime	UPTIME	AIX_LPAR	UNIXLPAR	Uptime	UPTIME
Virt_Context_CPU_Switches _per_Sec	VCCSPS	AIX_LPAR	UNIXLPAR	Virt_Context_CPU_Switches _per_Sec	VCCSPS

Table 30. Network_Adapters_Rates attribute group (table name: KPX_NETWORK_ADAPTERS_RATES)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Bandwidth_Util_Pct	BUP	Network	UNIXNET	Bandwidth_Util_Pct	BUP

Table 31. Paging_Space attribute group (table name: KPX_PAGING_SPACE)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Paging_Space_Free_Pct	FREE _PCT	Unix _Memory	UNIXMEM	Paging_Space_Free_Pct	FREE _PCT
Paging_Space_Used_Pct	USED _PCT	Unix _Memory	UNIXMEM	Paging_Space_Used_Pct	USED _PCT

Table 32. Physical_Memory attribute group (table name: KPX_PHYSICAL_MEMORY)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Comp_Memory	CM	Unix _Memory	UNIXMEM	Comp_Memory	CM
Decay_Rate	DECAY _RATE	Unix _Memory	UNIXMEM	Decay_Rate	DECAY _RATE
Non_Comp_Memory	NCM	Unix _Memory	UNIXMEM	Non_Comp_Memory	NCM
Repaging_Rate	RR	Unix _Memory	UNIXMEM	Repaging_Rate	RR

Table 33. Processes_Detail attribute group (table name: KPX_PROCESSES_DETAIL)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Page_Space_Used	PSU	Process	UNIXPS	Page_Space_Used	PSU
Resident_Data_Size	RDS	Process	UNIXPS	Resident_Data_Size	RDS
Resident_Text_Size	RTS	Process	UNIXPS	Resident_Text_Size	RTS
WLM_Name	WLM NAME	Process	UNIXPS	WLM_Name	WLM _NAME
WPAR_Name	WPAR _NAME	Process	UNIXPS	WPAR_Name	WPAR _NAME

Table 34. Virtual_Memory_Management attribute group (table name: KPX_VIRTUAL_MEMORY_MANAGEMENT)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Pages_Read_per_Sec	PRPS	Unix _Memory	UNIXMEM	Pages_Read_per_Sec	PRPS
Pages_Written_per_Sec	PWPS	Unix _Memory	UNIXMEM	Pages_Written_per_Sec	PWPS
Paging_Space_Read _per_Sec	PSRPS	Unix _Memory	UNIXMEM	Paging_Space_Read _per_Sec	PSRPS
Paging_Space_Write _per_Sec	PSWPS	Unix _Memory	UNIXMEM	Paging_Space_Write _per_Sec	PSWPS

Table 35. WPAR_CPU attribute group (table name: KPX_WPAR_CPU)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
CPU_Consumption_Limit	CCL	AIX _WPAR _CPU	UNIXWPARCP	CPU_Consumption_Limit	CCL
LPAR_CPU _Consumed_Pct	LCCP	AIX _WPAR _CPU	UNIXWPARCP	LPAR_CPU _Consumed_Pct	LCCP
LPAR_Entitlement	LE	AIX _WPAR _CPU	UNIXWPARCP	LPAR_Entitlement	LE
Num_CPUs_Consumed	NCC	AIX _WPAR _CPU	UNIXWPARCP	Num_CPUs_Consumed	NCC
RC_CPU_Limits _Hard_Max	RCLHM	AIX _WPAR _CPU	UNIXWPARCP	RC_CPU_Limits _Hard_Max	RCLHM

Table 35. WPAR_CPU attribute group (table name: KPX_WPAR_CPU) (continued)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
System_CPU_Pct	SCP	AIX _WPAR _CPU	UNIXWPARCP	System_CPU_Pct	SCP
System_Name	NODE	AIX _WPAR _CPU	UNIXWPARCP	System_Name	ORIGIN NODE
Timestamp	TIME STAMP	AIX _WPAR _CPU	UNIXWPARCP	Timestamp	TIME STAMP
User_CPU_Pct	UCP	AIX _WPAR _CPU	UNIXWPARCP	User_CPU_Pct	UCP
WPAR_CPU _Consumed_Pct	WCCP	AIX _WPAR _CPU	UNIXWPARCP	WPAR_CPU _Consumed_Pct	WCCP
WPAR_Name	WPAR _NAME	AIX _WPAR _CPU	UNIXWPARCP	WPAR_Name	WPAR _NAME

Table 36. WPAR_FileSystem attribute group (table name: KPX_WPAR_FILESYSTEM)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Device_Name	DN	AIX_WPAR _FileSystem	UNIXWPARFS	Device_Name	DN
Mount_Options	MO	AIX_WPAR _FileSystem	UNIXWPARFS	Mount_Options	MO
Mount_Point	MP	AIX_WPAR _FileSystem	UNIXWPARFS	Mount_Point	MP
Node_Name	NODE _NAME	AIX_WPAR _FileSystem	UNIXWPARFS	Node_Name	NODE _NAME
System_Name	NODE	AIX_WPAR _FileSystem	UNIXWPARFS	System_Name	ORIGIN NODE
Timestamp	TIMESTAMP	AIX_WPAR _FileSystem	UNIXWPARFS	Timestamp	TIMESTAMP
VFS_Type	VFS_TYPE	AIX_WPAR _FileSystem	UNIXWPARFS	VFS_Type	VFS_TYPE
WPAR_Name	WPAR _NAME	AIX_WPAR _FileSystem	UNIXWPARFS	WPAR_Name	WPAR _NAME

Table 37. WPAR_Information attribute group (table name: KPX_WPAR_INFORMATION)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Admin_Operation	AO	AIX_WPAR _Information	UNIXWPARIN	Admin_Operation	AO
Admin_Process_ID	API	AIX_WPAR _Information	UNIXWPARIN	Admin_Process_ID	API
Admin_Start_Time	AST	AIX_WPAR _Information	UNIXWPARIN	Admin_Start_Time	AST
Autostart	AUTOSTART	AIX_WPAR _Information	UNIXWPARIN	Autostart	AUTOSTART
Checkpointable	C	AIX_WPAR _Information	UNIXWPARIN	Checkpointable	C
Home	HOME	AIX_WPAR _Information	UNIXWPARIN	Home	HOME
Hostname	HOSTNAME	AIX_WPAR _Information	UNIXWPARIN	Hostname	HOSTNAME
IP_Address	IP_ADDRESS	AIX_WPAR _Information	UNIXWPARIN	IP_Address	IP_ADDRESS
Owner	OWNER	AIX_WPAR _Information	UNIXWPARIN	Owner	OWNER
RC_CPU_Limits _Hard_Max	RCLHM	AIX_WPAR _Information	UNIXWPARIN	RC_CPU_Limits _Hard_Max	RCLHM
RC_CPU _Limits_Min	RCLM	AIX_WPAR _Information	UNIXWPARIN	RC_CPU _Limits_Min	RCLM
RC_CPU _Limits_Soft_Max	RCLSM	AIX_WPAR _Information	UNIXWPARIN	RC_CPU _Limits_Soft_Max	RCLSM
RC_CPU_Shares	RCS	AIX_WPAR _Information	UNIXWPARIN	RC_CPU_Shares	RCS
RC_Is_Active	RIA	AIX_WPAR _Information	UNIXWPARIN	RC_Is_Active	RIA
RC_Max_Processes	RMP	AIX_WPAR _Information	UNIXWPARIN	RC_Max_Processes	RMP
RC_Max_Threads	RMT	AIX_WPAR _Information	UNIXWPARIN	RC_Max_Threads	RMT
RC_Memory _Limits_Hard_Max	RMLHM	AIX_WPAR _Information	UNIXWPARIN	RC_Memory _Limits_Hard_Max	RMLHM
RC_Memory _Limits_Min	RMLM	AIX_WPAR _Information	UNIXWPARIN	RC_Memory _Limits_Min	RMLM
RC_Memory _Limits_Soft_Max	RMLSM	AIX_WPAR _Information	UNIXWPARIN	RC_Memory _Limits_Soft_Max	RMLSM
RC_Memory_Shares	RMS	AIX_WPAR _Information	UNIXWPARIN	RC_Memory_Shares	RMS
RC_per_Process _VM_Limit	RPPVL	AIX_WPAR _Information	UNIXWPARIN	RC_per_Process _VM_Limit	RPPVL
RC_RSet	RC_RSET	AIX_WPAR _Information	UNIXWPARIN	RC_RSet	RC_RSET

Table 37. WPAR_Information attribute group (table name: KPX_WPAR_INFORMATION) (continued)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Shares_usr_Dir	SUD	AIX_WPAR _Information	UNIXWPARIN	Shares_usr_Dir	SUD
State	STATE	AIX_WPAR _Information	UNIXWPARIN	State	STATE
System_Name	NODE	AIX _Information	UNIXWPARIN	System_Name	ORIGIN NODE
Timestamp	TIMESTAMP	AIX _Information	UNIXWPARIN	Timestamp	TIMESTAMP
Type	TYPE	AIX_WPAR _Information	UNIXWPARIN	Type	TYPE
WPAR_Application _Path	WAP	AIX_WPAR _Information	UNIXWPARIN	WPAR_Application _Path	WAP
WPAR_Name	WPAR _NAME	AIX_WPAR _Information	UNIXWPARIN	WPAR_Name	WPAR _NAME

Table 38. WPAR_Network attribute group (table name: KPX_WPAR_NETWORK)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Broadcast_IP	BI	AIX_WPAR _Network	UNIXWPARNE	Broadcast_IP	BI
Interface_Name	IN	AIX_WPAR _Network	UNIXWPARNE	Interface_Name	IN
IP_Address	IP_ADDRESS	AIX_WPAR _Network	UNIXWPARNE	IP_Address	IP_ADDRESS
Network_Mask	NM	AIX_WPAR _Network	UNIXWPARNE	Network_Mask	NM
System_Name	NODE	AIX_WPAR _Network	UNIXWPARNE	System_Name	ORIGIN NODE
Timestamp	TIMESTAMP	AIX_WPAR _Network	UNIXWPARNE	Timestamp	TIMESTAMP
WPAR_Name	WPAR _NAME	AIX_WPAR _Network	UNIXWPARNE	WPAR_Name	WPAR _NAME

Table 39. WPAR_Physical_Memory attribute group (table name: KPX_WPAR_PHYSICAL_MEMORY)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Free_Memory_MB	FMM	AIX _WPAR _Physical _Memory	UNIXWPARPM	Free_Memory_MB	FMM

Table 39. WPAR_Physical_Memory attribute group (table name: KPX_WPAR_PHYSICAL_MEMORY) (continued)

AIX Premium agent attribute name	AIX Premiun agent column name	UNIX OS agent attribute group	UNIX OS agent table name	UNIX OS agent attribute name	UNIX OS agent column name
Free_Memory_Pct	FMP	AIX _WPAR _Physical _Memory	UNIXWPARPM	Free_Memory_Pct	FMP
LPAR_Memory_Size_MB	LMSM	AIX _WPAR _Physical _Memory	UNIXWPARPM	LPAR_Memory_Size_MB	LMSM
LPAR_Memory_Used_Pct	LMUP	AIX _WPAR _Physical _Memory	UNIXWPARPM	LPAR_Memory_Used_Pct	LMUP
Memory_Size_MB	MSM	AIX _WPAR _Physical _Memory	UNIXWPARPM	Memory_Size_MB	MSM
RC_Memory_Limits _Hard_Max	RMLHM	AIX _WPAR _Physical _Memory	UNIXWPARPM	RC_Memory_Limits _Hard_Max	RMLHM
System_Name	NODE	AIX _WPAR _Physical _Memory	UNIXWPARPM	System_Name	ORIGIN NODE
Timestamp	TIME STAMP	AIX _WPAR _Physical _Memory	UNIXWPARPM	Timestamp	TIME STAMP
Used_Memory_MB	UMM	AIX _WPAR _Physical _Memory	UNIXWPARPM	Used_Memory_MB	UMM
Used_Memory_Pct	UMP	AIX _WPAR _Physical _Memory	UNIXWPARPM	Used_Memory_Pct	UMP
WPAR_Name	WPAR _NAME	AIX _WPAR _Physical _Memory	UNIXWPARPM	WPAR_Name	WPAR _NAME

Appendix C. IBM Tivoli Enterprise Console event mapping

Specific event mapping is provided for those monitoring agents that support Distributed Monitoring migration. The specific event mapping creates Distributed Monitoring events for Distributed Monitoring migrated situations. For a list of these situations and their related event classes, see Table 40.

Generic event mapping provides useful event class and attribute information for situations that do not have specific event mapping defined. Each event class corresponds to an attribute group in the monitoring agent. For a description of the event slots for each event class, see Table 41 on page 248. For more information about mapping attribute groups to event classes, see the *IBM Tivoli Monitoring Administrator's Guide*.

BAROC files are found on the Tivoli Enterprise Monitoring Server in the installation directory in TECLIB (that is, *install_dir*/cms/TECLIB for Windows systems and *install_dir*/tables/*TEMS_hostname*/TECLIB for UNIX systems). For information on the current version of the BAROC file, see the *IBM Tivoli Monitoring Installation and Setup Guide*. IBM Tivoli Enterprise Console event synchronization provides a collection of ready-to-use rule sets that you can deploy with minimal configuration. Be sure to install IBM Tivoli Enterprise Console event synchronization to access the correct Sentry.baroc, which is automatically included during base configuration of IBM Tivoli Enterprise Console rules if you indicate that you want to use an existing rulebase. See the *IBM Tivoli Monitoring Installation and Setup Guide* for details.

Table 40. Overview of Distributed Monitoring migrated situations

Situation	IBM Tivoli Enterprise Console event class
UX_USInodes*	Sentry2_0_inodes Sentry2_0_inodesused
UX_USIUsPct*	Sentry2_0_inodesusedpct
UX_USDkUPct*	Sentry2_0_diskusedpct
UX_USDskAva*	Sentry2_0_diskavail
UX_USDskUsd*	Sentry2_0_diskused
UX_USDIORtK*	Sentry2_0_diskioratek
UX_USPDskRt*	Sentry2_0_peakdiskrate
UX_USPkDkXf*	Sentry2_0_peakdiskxfer
UX_USSpcUtl*	Sentry2_0_spaceutil
UX_USSpcUtK*	Sentry2_0_spaceutilkb
UX_USReqWt*	Sentry2_0_reqwait
UX_USReqTm*	Sentry2_0_reqtime
UX_USRPCTmO*	Sentry2_0_rpctmout
UX_USBadNFS*	Sentry2_0_badnfs
UX_USBadRPC*	Sentry2_0_badrpc
UX_USNtInEr*	Sentry2_0_netinerr
UX_USNtInEX*	Sentry2_0_netinerrx
UX_USNetIn*	Sentry2_0_netinerr

Table 40. Overview of Distributed Monitoring migrated situations (continued)

Situation	IBM Tivoli Enterprise Console event class
UX_USNetInX*	Sentry2_0_netinx
UX_USNetCol*	Sentry2_0_netcoll
UX_USNetCoX*	Sentry2_0_netcollx
UX_USNtCPct*	Sentry2_0_netcollpct
UX_USNCPctX*	Sentry2_0_netcollpctx
UX_USNetOEr*	Sentry2_0_netouterr
UX_USNetOEX*	Sentry2_0_netouterrx
UX_USNetOut*	Sentry2_0_netouterr
UX_USNetOX*	Sentry2_0_netoutx
UX_USNtCIRt*	Sentry2_0_netcollirate
UX_USNtIERt*	Sentry2_0_netinerrate
UX_USNtOERt*	Sentry2_0_netouterrate
UX_USNetIRt*	Sentry2_0_netinrate
UX_USNetORt*	Sentry2_0_netoutrate
UX_USSwpAva*	Sentry2_0_swapavail
UX_USTProcs*	Sentry2_0_totalprocs
UX_USCPUIdl*	Sentry2_0_cpuidle
UX_USCPUSys*	Sentry2_0_cpusys
UX_USCPUUsr*	Sentry2_0_cpuusr
UX_USCPUSpu*	Sentry2_0_cpuspu
UX_USZombie*	Sentry2_0_zombies
UX_USLdAv15*	Sentry2_0_loadavgfifteenm
UX_USLdAv5*	Sentry2_0_loadavgonem
UX_USLdAv1*	Sentry2_0_loadavgonem
UX_USPgScnR*	Sentry2_0_pagescanrate
UX_USPgIns*	Sentry2_0_pageins
UX_USPgOuts*	Sentry2_0_pageouts
UX_USPgScan*	Sentry2_0_pagescans
UX_USPgInRt*	Sentry2_0_pageinrate
UX_USPgORt*	Sentry2_0_pageoutrate
UX_USPgScRt*	Sentry2_0_pagescanrate
UX_USRnQJbs*	Sentry2_0_runqjobs
UX_USACPUBu*	Sentry2_0_avgcpubusy
UX_USACPUSy*	Sentry2_0_avgcpusys
UX_USACPUUs*	Sentry2_0_avgcpuusr
UX_USFilPrm*	Sentry2_0_fileperm
UX_USULginT*	Sentry2_0_ulogintot
UX_UDskAva*	universal_diskavail
UX_UDskUsd*	universal_diskused
UX_UDskUPct*	universal_diskusedpct

Table 40. Overview of Distributed Monitoring migrated situations (continued)

Situation	IBM Tivoli Enterprise Console event class
UX_UIndsFre*	universal_diskusedpct
UX_UINdsUsd*	universal_diskusedpct
UX_ULoadAvg*	universal_loadavg
UX_UPageOut*	universal_pageouts
UX_USwapAva*	universal_swapavail

To determine what event class is sent when a given situation is triggered, look at the first referenced attribute group in the situation predicate. The event class that is associated with that attribute group is the one that is sent. This is true for both pre-packaged situations and user-defined situations. See the table below for attribute group to event classes and slots mapping information.

For example, if the situation is monitoring the User Sessions attribute from the UNIX All Users Group attribute group, the event class that is sent once the situation is triggered is ITM_Unix_All_Users.

Note: There are cases where these mappings generate events that are too large for the Tivoli Enterprise Console. In these cases, the event class names and the event slot names are the same, but some of the event slots are omitted.

Each of the event classes is a child of KUX_Base. The KUX_Base event class can be used for generic rules processing for any event from the Monitoring Agent for UNIX OS.

Table 41. Overview of attribute groups to event classes and slots

Attribute group	event classes and slots
System	ITM_System event class with these slots: • system_name: STRING • timestamp: STRING • type: STRING • version: STRING • total_real_memory: INTEGER • total_real_memory_enum: STRING • total_virtual_memory: INTEGER • total_virtual_memory_enum: STRING • up_time: INTEGER • up_time_enum: STRING • users_session_number: INTEGER • system_procs_number: INTEGER • net_address: STRING • net_address_enum: STRING • user_cpu: INTEGER • user_cpu_enum: STRING • users_session_number_enum: STRING • system_cpu: INTEGER • system_cpu_enum: STRING • system_procs_number_enum: STRING • idle_cpu: INTEGER • idle_cpu_enum: STRING • wait_io: INTEGER • wait_io_enum: STRING • processes_in_run_queue: INTEGER • processes_in_run_queue_enum: STRING • processes_waiting: INTEGER • processes_waiting_enum: STRING • page_faults: INTEGER • page_faults_enum: STRING • page_reclaims: INTEGER • page_reclaims_enum: STRING • pages_paged_in: INTEGER • pages_paged_in_enum: STRING • pages_paged_out: INTEGER • pages_paged_out_enum: STRING • page_ins: INTEGER • page_ins_enum: STRING • page_outs: INTEGER • page_outs_enum: STRING • free_memory: INTEGER • free_memory_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
System (Cont.)	• active_virtual_memory: INTEGER • active_virtual_memory_enum: STRING • cpu_context_switches: INTEGER • cpu_context_switches_enum: STRING • system_calls: INTEGER • system_calls_enum: STRING • forks_executed: INTEGER • forks_executed_enum: STRING • execs_executed: INTEGER • execs_executed_enum: STRING • block_reads: INTEGER • block_reads_enum: STRING • block_writes: INTEGER • block_writes_enum: STRING • logical_block_reads: INTEGER • logical_block_reads_enum: STRING • logical_block_writes: INTEGER • logical_block_writes_enum: STRING • nonblock_reads: INTEGER • nonblock_reads_enum: STRING • nonblock_writes: INTEGER • nonblock_writes_enum: STRING • receive_interrupts: INTEGER • receive_interrupts_enum: STRING • transmit_interrupts: INTEGER • transmit_interrupts_enum: STRING • modem_interrupts: INTEGER • modem_interrupts_enum: STRING • active_internet_connections: INTEGER • active_internet_connections_enum: STRING • active_sockets: INTEGER • active_sockets_enum: STRING • load_average_1_min: REAL • load_average_1_min_enum: STRING • load_average_5_min: REAL • load_average_5_min_enum: STRING • load_average_15_min: REAL • load_average_15_min_enum: STRING • dummy-memory_free: INTEGER • dummy-memory_free_enum: STRING • memory_used: INTEGER • memory_used_enum: STRING • page_scan_rate: INTEGER • page_scan_rate_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
System (Cont.)	• virtual_memory_percent_used: REAL • virtual_memory_percent_used_enum: STRING • virtual_memory_percent_available: REAL • virtual_memory_percent_available_enum: STRING • cpu_busy: INTEGER • cpu_busy_enum: STRING • system_read: INTEGER • system_read_enum: STRING • system_write: INTEGER • system_write_enum: STRING • system_threads: INTEGER • system_threads_enum: STRING • processes_runnable: INTEGER • processes_runnable_enum: STRING • processes_running: INTEGER • processes_running_enum: STRING • processes_sleeping: INTEGER • processes_sleeping_enum: STRING • processes_idle: INTEGER • processes_idle_enum: STRING • processes_zombie: INTEGER • processes_zombie_enum: STRING • processes_stopped: INTEGER • processes_stopped_enum: STRING • threads_in_run_queue: INTEGER • threads_in_run_queue_enum: STRING • threads_waiting: INTEGER • threads_waiting_enum: STRING • boot_time: STRING • pending_io_waits: INTEGER • pending_io_waits_enum: STRING • start_io: INTEGER • start_io_enum: STRING • device_interrupts: INTEGER • device_interrupts_enum: STRING • uptime: STRING • parameter: STRING • omunx_value: STRING • swap_space_free: INTEGER • swap_space_free_enum: STRING • page_ins_rate: INTEGER • page_ins_rate_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
System (Cont.)	• page_out_rate: INTEGER • page_out_rate_enum: STRING • page_scanning: INTEGER • page_scanning_enum: STRING • avg_pageins_1: INTEGER • avg_pageins_1_enum: STRING • avg_pageins_5: INTEGER • avg_pageins_5_enum: STRING • avg_pageins_15: INTEGER • avg_pageins_15_enum: STRING • avg_pageins_60: INTEGER • avg_pageins_60_enum: STRING • avg_pageout_1: INTEGER • avg_pageout_1_enum: STRING • avg_pageout_5: INTEGER • avg_pageout_5_enum: STRING • avg_pageout_15: INTEGER • avg_pageout_15_enum: STRING • avg_pageout_60: INTEGER • avg_pageout_60_enum: STRING • avg_pagescan_1: INTEGER • avg_pagescan_1_enum: STRING • avg_pagescan_5: INTEGER • avg_pagescan_5_enum: STRING • avg_pagescan_15: INTEGER • avg_pagescan_15_enum: STRING • avg_pagescan_60: INTEGER • avg_pagescan_60_enum: STRING • avg_proccesses_runqueue_60: INTEGER • avg_proccesses_runqueue_60_enum: STRING • ipv6_address: STRING • ipv6_address_enum: STRING • zone_id: INTEGER • zone_id_enum: STRING • zone_name: STRING • zone_name_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Disk	ITM_Disk event class with these slots: • system_name: STRING • timestamp: STRING • name: STRING • mount_point: STRING • size: INTEGER • size_enum: STRING • space_used: INTEGER • space_used_enum: STRING • space_available: INTEGER • space_available_enum: STRING • inode_size: INTEGER • inode_size_enum: STRING • inodes_used: INTEGER • inodes_used_enum: STRING • inodes_free: INTEGER • inodes_free_enum: STRING • space_used_percent: INTEGER • space_used_percent_enum: STRING • inodes_used_percent: INTEGER • inodes_used_percent_enum: STRING • fs_type: STRING • space_available_percent: INTEGER • space_available_percent_enum: STRING • name_u: STRING • mount_point_u: STRING • size_mb: INTEGER • size_mb_enum: STRING • size_gb: INTEGER • size_gb_enum: STRING • space_used_mb: INTEGER • space_used_mb_enum: STRING • space_used_gb: INTEGER • space_used_gb_enum: STRING • space_available_mb: INTEGER • space_available_mb_enum: STRING • space_available_gb: INTEGER • space_available_gb_enum: STRING • inodes_available_percent: INTEGER • inodes_available_percent_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Disk (Continued)	• size_64: REAL • size_64_enum: STRING • space_used_64: REAL • space_used_64_enum: STRING • space_available_64: REAL • space_available_64_enum: STRING • inode_size_64: REAL • inode_size_64_enum: STRING • inodes_used_64: REAL • inodes_used_64_enum: STRING • inodes_free_64: REAL • inodes_free_64_enum: STRING • size_mb_decimal: REAL • size_mb_decimal_enum: STRING • size_gb_decimal: REAL • size_gb_decimal_enum: STRING • space_used_mb_decimal: REAL • space_used_mb_decimal_enum: STRING • space_used_gb_decimal: REAL • space_used_gb_decimal_enum: STRING • space_available_mb_decimal: REAL • space_available_mb_decimal_enum: STRING • space_available_gb_decimal: REAL • space_available_gb_decimal_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Disk_Performance	ITM_Disk_Performance event class with these slots: • system_name: STRING • timestamp: STRING • disk_name: STRING • transfer_rate: INTEGER • transfer_rate_enum: STRING • transferred_bytes: INTEGER • transferred_bytes_enum: STRING • busy_percent: INTEGER • busy_percent_enum: STRING • avg_queue: INTEGER • avg_queue_enum: STRING • avg_wait: INTEGER • avg_wait_enum: STRING • avg_serv: INTEGER • avg_serv_enum: STRING • disk_name_u: STRING • percent_disk_read_time: INTEGER • percent_disk_read_time_enum: STRING • percent_disk_write_time: INTEGER • percent_disk_write_time_enum: STRING • disk_reads_sec: INTEGER • disk_reads_sec_enum: STRING • disk_writes_sec: INTEGER • disk_writes_sec_enum: STRING • disk_read_bytes_sec: INTEGER • disk_read_bytes_sec_enum: STRING • disk_write_bytes_sec: INTEGER • disk_write_bytes_sec_enum: STRING • avg_disk_bytes_xfer: INTEGER • avg_disk_bytes_xfer_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Network	ITM_Network event class with these slots: • system_name: STRING • timestamp: STRING • network_interface_name: STRING • network_interface_name_enum: STRING • interface_ip_address: STRING • interface_dns_name: STRING • interface_dns_name_enum: STRING • interface_status: STRING • interface_status_enum: STRING • transmission_unit_maximum: INTEGER • transmission_unit_maximum_enum: STRING • received_count: INTEGER • received_count_enum: STRING • transmitted_count: INTEGER • transmitted_count_enum: STRING • frames_received: INTEGER • frames_received_enum: STRING • frames_transmitted: INTEGER • frames_transmitted_enum: STRING • input_errors: INTEGER • input_errors_enum: STRING • output_errors: INTEGER • output_errors_enum: STRING • collisions: INTEGER • collisions_enum: STRING • subunit_driver: INTEGER • subunit_driver_enum: STRING • avg_coll_rate_1: INTEGER • avg_coll_rate_1_enum: STRING • avg_coll_rate_5: INTEGER • avg_coll_rate_5_enum: STRING • avg_coll_rate_15: INTEGER • avg_coll_rate_15_enum: STRING • avg_coll_rate_60: INTEGER • avg_coll_rate_60_enum: STRING • avg_in_rate_1: INTEGER • avg_in_rate_1_enum: STRING • avg_in_rate_5: INTEGER • avg_in_rate_5_enum: STRING • avg_in_rate_15: INTEGER • avg_in_rate_15_enum: STRING • avg_in_rate_60: INTEGER • avg_in_rate_60_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Network (continued)	• avg_inerr_rate_1: INTEGER • avg_inerr_rate_1_enum: STRING • avg_inerr_rate_5: INTEGER • avg_inerr_rate_5_enum: STRING • avg_inerr_rate_15: INTEGER • avg_inerr_rate_15_enum: STRING • avg_inerr_rate_60: INTEGER • avg_inerr_rate_60_enum: STRING • avg_out_rate_1: INTEGER • avg_out_rate_1_enum: STRING • avg_out_rate_5: INTEGER • avg_out_rate_5_enum: STRING • avg_out_rate_15: INTEGER • avg_out_rate_15_enum: STRING • avg_out_rate_60: INTEGER • avg_out_rate_60_enum: STRING • avg_outerr_rate_1: INTEGER • avg_outerr_rate_1_enum: STRING • avg_outerr_rate_5: INTEGER • avg_outerr_rate_5_enum: STRING • avg_outerr_rate_15: INTEGER • avg_outerr_rate_15_enum: STRING • avg_outerr_rate_60: INTEGER • avg_outerr_rate_60_enum: STRING • received_mb: REAL • received_mb_enum: STRING • transmitted_mb: REAL • transmitted_mb_enum: STRING • received_mb_total: REAL • received_mb_total_enum: STRING • transmitted_mb_total: REAL • transmitted_mb_total_enum: STRING • ipv4_dns_name: STRING • ipv4_dns_name_enum: STRING • type: INTEGER • type_enum: STRING • input_packet_errors_percent: INTEGER • input_packet_errors_percent_enum: STRING • output_packet_errors_percent: INTEGER • output_packet_errors_percent_enum: STRING • packet_collisions_percent: INTEGER • packet_collisions_percent_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Network (continued)	• mac_address: STRING • mac_address_enum: STRING • received_count_64: REAL • received_count_64_enum: STRING • transmitted_count_64: REAL • transmitted_count_64_enum: STRING
User	ITM_User event class with these slots: • system_name: STRING • timestamp: STRING • login_name: STRING • name: STRING • terminal: STRING • idle_time: INTEGER • idle_time_enum: STRING • login_time: STRING • location: STRING • user_id: INTEGER • user_id_enum: STRING • process_id: INTEGER • login_name_u: STRING • name_u: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Process	ITM_Process event class with these slots: • system_name: STRING • timestamp: STRING • process_id: INTEGER • process_id_enum: STRING • flag: STRING • execution_state: STRING • execution_state_enum: STRING • user_id: INTEGER • user_id_enum: STRING • parent_process_id: INTEGER • parent_process_id_enum: STRING • cpu_utilization: INTEGER • cpu_utilization_enum: STRING • priority: INTEGER • priority_enum: STRING • nice_value: INTEGER • nice_value_enum: STRING • entry_address: STRING • size: INTEGER • size_enum: STRING • event_waited_on: STRING • terminal_device: STRING • time: STRING • command: STRING • process_command: STRING • reptype: STRING • real_group_id: INTEGER • real_group_id_enum: STRING • effective_user_id: INTEGER • effective_user_id_enum: STRING • effective_group_id: INTEGER • effective_group_id_enum: STRING • process_group_leader_id: INTEGER • process_group_leader_id_enum: STRING • session_id: INTEGER • session_id_enum: STRING • scheduling_class: STRING • cpu_id: INTEGER • cpu_id_enum: STRING • user_name: STRING • starttime: STRING • elapsed_time: STRING • virtual_size: INTEGER • virtual_size_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Process (continued)	• mem_pct: REAL • mem_pct_enum: STRING • cpu_pct: REAL • cpu_pct_enum: STRING • total_cpu_percent: REAL • total_cpu_percent_enum: STRING • sample_cpu_pct: REAL • sample_cpu_pct_enum: STRING • heap_size: INTEGER • heap_size_enum: STRING • stack_size: INTEGER • stack_size_enum: STRING • major_fault: INTEGER • major_fault_enum: STRING • minor_fault: INTEGER • minor_fault_enum: STRING • context_switch: INTEGER • context_switch_enum: STRING • involuntary_context_switch: INTEGER • involuntary_context_switch_enum: STRING • user_cpu_time: STRING • system_cpu_time: STRING • total_cpu_time: STRING • thread_count: INTEGER • thread_count_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Process (Cont.)	• child_user_cpu_time: STRING • child_user_cpu_time_enum: STRING • child_system_cpu_time: STRING • child_system_cpu_time_enum: STRING • total_child_cpu_time: STRING • total_child_cpu_time_enum: STRING • wait_cpu_time: STRING • wait_cpu_time_enum: STRING • wait_lock_time: STRING • wait_lock_time_enum: STRING • read_per_write: INTEGER • read_per_write_enum: STRING • cpu_time: INTEGER • cpu_time_enum: STRING • parameter: STRING • omunx_value: STRING • command_u: STRING • process_command_u: STRING • user_name_u: STRING • group_name: STRING • effective_user_name: STRING • effective_group_name: STRING • base_command: STRING • type: STRING • zone_id: INTEGER • zone_id_enum: STRING • zone_name: STRING • zone_name_enum: STRING • process_count: INTEGER • process_count_enum: STRING • major_fault_64: REAL • major_fault_64_enum: STRING • minor_fault_64: REAL • minor_fault_64_enum: STRING • context_switch_64: REAL • context_switch_64_enum: STRING • involuntary_context_switch_64: REAL • involuntary_context_switch_64_enum: STRING • read_per_write_64: REAL • read_per_write_64_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
File_Information	ITM_File_Information event class with these slots: • system_name: STRING • timestamp: STRING • path: STRING • file: STRING • size: INTEGER • size_enum: STRING • owner: STRING • group: STRING • last_changed_time: STRING • last_accessed_time: STRING • links: INTEGER • links_enum: STRING • access: INTEGER • access_enum: STRING • type: STRING • type_enum: STRING • link_name: STRING • path_u: STRING • file_u: STRING • owner_u: STRING • group_u: STRING • link_name_u: STRING • size_mb: INTEGER • size_mb_enum: STRING • mode: STRING • last_attr_chg_time: STRING • checksum: STRING • checksum_algorithm: INTEGER • checksum_algorithm_enum: STRING • file_content_changed: INTEGER • file_content_changed_enum: STRING • size_64: REAL • size_64_enum: STRING • size_mb_64: REAL • size_mb_64_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
N_F_S_and_R_P_C_Statistics	ITM_N_F_S_and_R_P_C_Statistics event class with these slots: • system_name: STRING • timestamp: STRING • rpc_server_calls_rejected: INTEGER • rpc_server_calls_rejected_enum: STRING • rpc_server_times_rpc_packet_unavailable: INTEGER • rpc_server_times_rpc_packet_unavailable_enum: STRING • rpc_server_packets_too_short: INTEGER • rpc_server_packets_too_short_enum: STRING • rpc_server_packets_with_malformed_header: INTEGER • rpc_server_packets_with_malformed_header_enum: STRING • rpc_client_calls_rejected_by_server: INTEGER • rpc_client_calls_rejected_by_server_enum: STRING • rpc_client_calls_retransmitted: INTEGER • rpc_client_calls_retransmitted_enum: STRING • rpc_client_replies_not_matching_calls: INTEGER • rpc_client_replies_not_matching_calls_enum: STRING • rpc_client_calls_timed_out: INTEGER • rpc_client_calls_timed_out_enum: STRING • rpc_client_times_call_wait_on_busy: INTEGER • rpc_client_times_call_wait_on_busy_enum: STRING • rpc_client_times_authentication_refreshed: INTEGER • rpc_client_times_authentication_refreshed_enum: STRING • nfs_server_calls: INTEGER • nfs_server_calls_enum: STRING • nfs_server_calls_rejected: INTEGER • nfs_server_calls_rejected_enum: STRING • nfs_server_rejected_call_percentage: INTEGER • nfs_server_rejected_call_percentage_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
N_F_S_and_R_P_C_Statistics (Continued)	• nfs_server_null_calls: INTEGER • nfs_server_null_calls_enum: STRING • nfs_server_get_attribute_calls: INTEGER • nfs_server_get_attribute_calls_enum: STRING • nfs_server_set_attribute_calls: INTEGER • nfs_server_set_attribute_calls_enum: STRING • nfs_server_root_calls: INTEGER • nfs_server_root_calls_enum: STRING • nfs_server_lookups: INTEGER • nfs_server_lookups_enum: STRING • nfs_server_read_link_calls: INTEGER • nfs_server_read_link_calls_enum: STRING • nfs_server_read_calls: INTEGER • nfs_server_read_calls_enum: STRING • nfs_server_write_cache_calls: INTEGER • nfs_server_write_cache_calls_enum: STRING • nfs_server_writes: INTEGER • nfs_server_writes_enum: STRING • nfs_server_file_creates: INTEGER • nfs_server_file_creates_enum: STRING • nfs_server_remove_file_calls: INTEGER • nfs_server_remove_file_calls_enum: STRING • nfs_server_rename_file_calls: INTEGER • nfs_server_rename_file_calls_enum: STRING • nfs_server_link_calls: INTEGER • nfs_server_link_calls_enum: STRING • nfs_server_symbolic_link_calls: INTEGER • nfs_server_symbolic_link_calls_enum: STRING • nfs_server_make_directory_calls: INTEGER • nfs_server_make_directory_calls_enum: STRING • nfs_server_remove_directory_calls: INTEGER • nfs_server_remove_directory_calls_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
N_F_S_and_R_P_C_Statistics (Continued)	• nfs_server_read_directory_calls: INTEGER • nfs_server_read_directory_calls_enum: STRING • nfs_server_file_system_statistics_calls: INTEGER • nfs_server_file_system_statistics_calls_ enum: STRING • nfs_client_calls: INTEGER • nfs_client_calls_enum: STRING • nfs_client_calls_rejected: INTEGER • nfs_client_calls_rejected_enum: STRING • nfs_client_rejected_call_percentage: INTEGER • nfs_client_null_calls: INTEGER • nfs_client_get_attribute_calls: INTEGER • nfs_client_set_attribute_calls: INTEGER • nfs_client_root_calls: INTEGER • nfs_client_lookups: INTEGER • nfs_client_read_link_calls: INTEGER • nfs_client_read_calls: INTEGER • nfs_client_write_cache_calls: INTEGER • nfs_client_writes: INTEGER • nfs_client_file_creates: INTEGER • nfs_client_remove_file_calls: INTEGER • nfs_client_rename_file_calls: INTEGER • nfs_client_link_calls: INTEGER • nfs_client_symbolic_link_calls: INTEGER • nfs_client_make_directory_calls: INTEGER • nfs_client_remove_directory_calls: INTEGER • nfs_client_read_directory_calls: INTEGER • nfs_client_file_system_statistics_calls: INTEGER • parameter: STRING • omunx_value: STRING • nfs_version: INTEGER • nfs_version_enum: STRING • rpc_server_calls: INTEGER • rpc_server_calls_enum: STRING • rpc_server_dup_checks: INTEGER • rpc_server_dup_checks_enum: STRING • rpc_server_dup_reqs: INTEGER • rpc_server_dup_reqs_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
N_F_S_and_R_P_C_Statistics (Continued)	• rpc_client_calls: INTEGER • rpc_client_calls_enum: STRING • rpc_server_dup_reqs_percent: INTEGER • rpc_server_dup_reqs_percent_enum: STRING • rpc_server_calls_rejected_percent: INTEGER • rpc_server_calls_rejected_percent_ enum: STRING • rpc_client_calls_retransmitted_percent: INTEGER • rpc_client_calls_retransmitted_percent_ enum: STRING • rpc_client_calls_retransmitted_limit_ percent: INTEGER • rpc_client_calls_retransmitted_limit_ percent_enum: STRING; • rpc_client_calls_timed_out_percent: INTEGER • rpc_client_calls_timed_out_percent_enum: STRING • rpc_client_bad_xid_replies_percent: INTEGER • rpc_client_bad_xid_replies_percent_enum: STRING; • rpc_client_bad_xid_replies_limit_percent: INTEGER • rpc_client_bad_xid_replies_limit_percent_ enum: STRING • rpc_client_calls_rejected_by_server_percent: INTEGER • rpc_client_calls_rejected_by_server_ percent_enum: STRING • nfs_server_get_attr_percent: INTEGER • nfs_server_get_attr_percent_enum: STRING • nfs_server_write_percent: INTEGER • nfs_server_write_percent_enum: STRING • nfs_server_read_percent: INTEGER • nfs_server_read_percent_enum: STRING • nfs_server_read_link_percent: INTEGER • nfs_server_read_link_percent_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
SMP_CPU	ITM_SMP_CPU event class with these slots: • system_name: STRING • timestamp: STRING • cpu_id: INTEGER • cpu_id_enum: STRING • user_cpu: INTEGER • user_cpu_enum: STRING • system_cpu: INTEGER • system_cpu_enum: STRING • idle_cpu: INTEGER • idle_cpu_enum: STRING • wait_io: INTEGER • wait_io_enum: STRING • cpu_busy: INTEGER • cpu_busy_enum: STRING • minor_faults: INTEGER • minor_faults_enum: STRING • major_faults: INTEGER • major_faults_enum: STRING • cross_calls: INTEGER • cross_calls_enum: STRING • interrupts: INTEGER • interrupts_enum: STRING • interrupts_as_threads: INTEGER • interrupts_as_threads_enum: STRING • context_switches: INTEGER • context_switches_enum: STRING • involuntary_context_switches: INTEGER • involuntary_context_switches_enum: STRING • thread_migrations: INTEGER • thread_migrations_enum: STRING • spins_on_mutexes: INTEGER • spins_on_mutexes_enum: STRING • spins_on_rw_locks: INTEGER • spins_on_rw_locks_enum: STRING • system_calls: INTEGER • system_calls_enum: STRING • cpu_status: INTEGER • cpu_status_enum: STRING • parameter: STRING • omunx_value: STRING • cpu_usage: INTEGER • cpu_usage_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
SMP_CPU (continued)	• cpu_time: INTEGER • cpu_time_enum: STRING • avg_cpu_busy_1: INTEGER • avg_cpu_busy_1_enum: STRING • avg_cpu_busy_5: INTEGER • avg_cpu_busy_5_enum: STRING • avg_cpu_busy_15: INTEGER • avg_cpu_busy_15_enum: STRING • avg_cpu_busy_60: INTEGER • avg_cpu_busy_60_enum: STRING • avg_cpu_sys_1: INTEGER • avg_cpu_sys_1_enum: STRING • avg_cpu_sys_5: INTEGER • avg_cpu_sys_5_enum: STRING • avg_cpu_sys_15: INTEGER • avg_cpu_sys_15_enum: STRING • avg_cpu_sys_60: INTEGER • avg_cpu_sys_60_enum: STRING • avg_cpu_usr_1: INTEGER • avg_cpu_usr_1_enum: STRING • avg_cpu_usr_5: INTEGER • avg_cpu_usr_5_enum: STRING • avg_cpu_usr_15: INTEGER • avg_cpu_usr_15_enum: STRING • avg_cpu_usr_60: INTEGER • avg_cpu_usr_60_enum: STRING
Print_Queue	ITM_Unix_Print_Queue event class with these slots: • system_name: STRING • timestamp: STRING • print_queue_name: STRING • device_name: STRING • print_queue_status: STRING • print_queue_depth: INTEGER • print_queue_depth_enum: STRING • print_queue_job_size: INTEGER • print_queue_job_size_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Group	ITM_Unix_Group event class with these slots: • system_name: STRING • timestamp: STRING • group_name: STRING • group_id: INTEGER • group_id_enum: STRING • group_duplicated: INTEGER • group_duplicated_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Memory	ITM_Unix_Memory event class with these slots: • system_name: STRING • timestamp: STRING • total_virtual_storage_mb: INTEGER • total_virtual_storage_mb_enum: STRING • used_virtual_storage_mb: INTEGER • used_virtual_storage_mb_enum: STRING • avail_virtual_storage_mb: INTEGER • avail_virtual_storage_mb_enum: STRING • virtual_storage_pct_used: REAL • virtual_storage_pct_used_enum: STRING • virtual_storage_pct_avail: REAL • virtual_storage_pct_avail_enum: STRING • total_swap_space_mb: INTEGER • total_swap_space_mb_enum: STRING • used_swap_space_mb: INTEGER • used_swap_space_mb_enum: STRING • avail_swap_space_mb: INTEGER • avail_swap_space_mb_enum: STRING • used_swap_space_pct: REAL • used_swap_space_pct_enum: STRING • avail_swap_space_pct: REAL • avail_swap_space_pct_enum: STRING • total_real_mem_mb: INTEGER • total_real_mem_mb_enum: STRING • used_real_mem_mb: INTEGER • used_real_mem_mb_enum: STRING • avail_real_mem_mb: INTEGER • avail_real_mem_mb_enum: STRING • used_real_mem_pct: REAL • used_real_mem_pct_enum: STRING • avail_real_mem_pct: REAL • avail_real_mem_pct_enum: STRING • page_faults: INTEGER • page_faults_enum: STRING • page_reclaims: INTEGER • page_reclaims_enum: STRING • page_ins: INTEGER • page_ins_enum: STRING • page_outs: INTEGER • page_outs_enum: STRING • page_in_reqs: INTEGER • page_in_reqs_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Memory (continued)	• page_out_reqs: INTEGER • page_out_reqs_enum: STRING • page_in_kb_s: INTEGER • page_in_kb_s_enum: STRING • page_out_kb_s: INTEGER • page_out_kb_s_enum: STRING • page_in_1min: INTEGER • page_in_1min_enum: STRING • page_in_5min: INTEGER • page_in_5min_enum: STRING • page_in_15min: INTEGER • page_in_15min_enum: STRING • page_in_60min: INTEGER • page_in_60min_enum: STRING • page_out_1min: INTEGER • page_out_1min_enum: STRING • page_out_5min: INTEGER • page_out_5min_enum: STRING • page_out_15min: INTEGER • page_out_15min_enum: STRING • page_out_60min: INTEGER • page_out_60min_enum: STRING • page_scan: INTEGER • page_scan_enum: STRING • page_scan_kb: INTEGER • page_scan_kb_enum: STRING • page_scan_1min: INTEGER • page_scan_1min_enum: STRING • page_scan_5min: INTEGER • page_scan_5min_enum: STRING • page_scan_15min: INTEGER • page_scan_15min_enum: STRING • page_scan_60min: INTEGER • page_scan_60min_enum: STRING
File_Pattern	ITM_Unix_File_Pattern event class with these slots: • system_name: STRING • timestamp: STRING • file_name: STRING • match_pattern: STRING • match_option: INTEGER • match_option_enum: STRING • match_count: INTEGER • match_count_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
File_Comparison	ITM_Unix_File_Comparison event class with these slots: • system_name: STRING • timestamp: STRING • file_name_1: STRING • file_name_2: STRING • file_compare_option: INTEGER • file_compare_option_enum: STRING • file_compare_result: INTEGER • file_compare_result_enum: STRING
Ping	ITM_Unix_Ping event class with these slots: • system_name: STRING • timestamp: STRING • target_host: STRING • ping_result: INTEGER • ping_result_enum: STRING • response_time: REAL • response_time_enum: STRING
All_Users	ITM_Unix_All_Users event class with these slots: • system_name: STRING • timestamp: STRING • name: STRING • user_id: INTEGER • user_id_enum: STRING • password_null: INTEGER • password_null_enum: STRING • user_duplicated: INTEGER • user_duplicated_enum: STRING • user_sessions: INTEGER • user_sessions_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Solaris_Zones	ITM_Solaris_Zones event class with these slots: • system_name: STRING • timestamp: STRING • name: STRING • name_enum: STRING • zone_id: INTEGER • zone_id_enum: STRING • omunx_status: INTEGER • omunx_status_enum: STRING • path: STRING • pool_id: INTEGER • pool_id_enum: STRING • init_pid: INTEGER • init_pid_enum: STRING • zone_cpu_usage: REAL • zone_cpu_usage_enum: STRING • physical_memory: INTEGER • physical_memory_enum: STRING • virtual_memory: INTEGER • virtual_memory_enum: STRING • total_cpus: INTEGER • total_cpus_enum: STRING • scheduler: STRING • scheduler_enum: STRING • cpu_shares: INTEGER • cpu_shares_enum: STRING • cpu_share_pct: REAL • cpu_share_pct_enum: STRING

Table 41. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Machine_Information	ITM_Machine_Information event class with these slots: • system_name: STRING • timestamp: STRING • omunx_hostname: STRING • omunx_hostname_enum: STRING • hardware_manufacturer: STRING • hardware_manufacturer_enum: STRING • hardware_model: STRING • hardware_model_enum: STRING • number_of_physical_processors: INTEGER • number_of_physical_processors_enum: STRING • processor_megahertz: INTEGER • processor_megahertz_enum: STRING • machine_serial: STRING • machine_serial_enum: STRING • system_board_uuid: STRING • system_board_uuid_enum: STRING • virtual_machine_identifier: STRING • virtual_machine_identifier_enum: STRING
IP_Address	ITM_IP_Address event class with these slots: • system_name: STRING • timestamp: STRING • network_interface_name: STRING • ip_address: STRING • dns_name: STRING • dns_name_enum: STRING • ip_version: INTEGER • ip_version_enum: STRING

Appendix D. Historical data

Historical reports use a column header associated with a shorter character name that identifies the attributes. The tables in this appendix identify the historical table, the Monitoring Agent for UNIX OS attribute group associated with the table, the historical table column head (in capital letters), and the associated attribute name. Use the information in this appendix in conjunction with the information in Chapter 4, "Attributes reference," on page 33 to identify attribute definitions for the historical data tables. For any attribute *XXX* accompanied by an *XXX*_WHSC attribute, the Summarization and Pruning Agent calculates averages and totals based on (*XXX* * *XXX*_WHSC).

UNIXALLUSR historical table

The UNIXALLUSR historical table corresponds to the All Users attributes.

Table 42 lists the historical table column heads alphabetically and the corresponding All Users group attributes.

Table 42. UNIXALLUSR table column heads and the corresponding All Users attributes

Historical table column head	Attribute name
ORIGINNODE	System_Name
PWNULL	Password_Null
TIMESTAMP	Timestamp
UID	User_ID
USERDUP	User_Duplicated
USERNAME	Name
USERSES	User_Sessions

UNIXAMS historical table

The UNIXAMS historical table corresponds to the AIX AMS attributes.

Table 43 lists the historical table column heads alphabetically and the corresponding AIX AMS group attributes.

Table 43. UNIXAMS table column heads and the corresponding AIX AMS attributes

Historical table column head	Attribute name
AME	AMS_Mem_Entitlement
AMEI	AMS_Mem_Ent_InUse
AML	AMS_Mem_Loaned
AMS_MODE	AMS_Mode
API	AMS_Pool_ID
APM	AMS_Physical_Mem
APS	AMS_Pool_Size
HPIT	Hypervisor_Page_Ins

Table 43. UNIXAMS table column heads and the corresponding AIX AMS attributes (continued)

Historical table column head	Attribute name
ORIGINNODE	System_Name
TIMESTAMP	Timestamp

UNIXCPU historical table

The UNIXCPU historical table corresponds to the SMP CPU attributes.

Table 44 lists the historical table column heads alphabetically and the corresponding SMP CPU attributes.

Table 44. UNIXCPU table column heads and the corresponding SMP CPU attributes

Historical table column head	Attribute name
CPUBUSY	CPU_Busy
CPUID	CPU_ID
CPUSTAT	CPU_Status
CSPS	Context_Switches_per_Sec
ICSW	Involuntary_Context_Switches
IDLECPU	Idle_CPU
INTRRUPT	Interrupts
INTRTHRD	Interrups_As_Threads
LCS	Logical_Context_Switches
MAJF	Major_Faults
MINF	Minor_Faults
ORIGINNODE	System_Name
PC	Physical_Consumption
SMTX	Spins_On_Mutexes
SRWLOCKS	Spins_On_RW_Locks
SYSCALL	System_Calls
SYSCPU	System_CPU
THRDMIGR	Thread_Migrations
TIMESTAMP	Timestamp
USRCPU	User_CPU
WAITIO	Wait_I/O
XCALLS	Cross_Calls
ZATTRIB	Parameter *Column seen in historical data collection tables but currently not collecting validated data.*
ZVALUE	Value *Column seen in historical data collection tables but currently not collecting validated data.*

UNIXDEVIC historical table

The UNIXDEVIC historical table corresponds to the AIX Devices attributes.

Table 45 lists the historical table column heads alphabetically and the corresponding AIX Devices group attributes.

Table 45. UNIXDEVIC table column heads and the corresponding AIX Devices attributes

Historical table column head	Attribute name
DCLASS	DClass
NAME	Name
ORIGINNODE	System_Name
PARENT	Parent
STATE	State
TIMESTAMP	Timestamp
TYPE	Type

UNIXDISK historical table

The UNIXDISK historical table corresponds to the Disk Information attributes.

Table 46 lists the historical table column heads alphabetically and the corresponding Disk group attributes.

Table 46. UNIXDISK table column heads and the corresponding Disk Information attributes

Historical table column head	Attribute name
DSKNAME	Name
DSKSIZE	Size
FSTYPE	FS_Type *Column seen in historical data collection tables but currently not collecting validated data.*
INODEFREE	Inodes_Free
INODESIZE	Inode_Size
INODEUSED	Inodes_Used
MOUNTPT	Mount_Point
ORIGINNODE	System_Name
PCTINDUSED	Inodes_Used_Percent
PCTSPCAV	Space_Available_Percent *Column seen in historical data collection tables but currently not collecting validated data.*
PCTSPCUSED	Space_Used_Percent
SPCAVAIL	Space_Available
SPCUSED	Space_Used
TIMESTAMP	Timestamp
VGN	Volume_Group_Name

UNIXDPERF historical table

The UNIXDPERF historical table corresponds to the Disk Performance attributes.

Table 47 lists the historical table column heads alphabetically and the corresponding Disk Performance attributes.

Table 47. UNIXDPERF table column heads and the corresponding Disk Performance attributes

Historical table column head	Attribute name
ASS	Avg_ServiceQ_Size
AVGSERV	Avg_Serv
DSKAVQUE	Avg_Queue
DSKAVWAIT	Avg_Wait
DSKBUSY	Busy_Percent
DSKBYTESIZ	Transferred_Bytes
DSKNAME	Disk_Name
DSKXFERRAT	Transfer_Rate
ORIGINNODE	System_Name
PARENT	Parent
SFPS	ServiceQ_Full_per_Sec
TIMESTAMP	Timestamp
TKPS	Transfers_KB_per_Sec
TYPE	Type

UNIXDUSERS historical table

The UNIXDUSERS historical table corresponds to the AIX Defined Users attributes.

Table 48 lists the historical table column heads alphabetically and the corresponding AIX Defined Users attributes.

Table 48. UNIXDUSERS table column heads and the corresponding AIX Defined Users attributes

Historical table column head	Attribute name
AL	Account_Locked
EXPIRES	Expires
L	Loginretries
ROLES	Roles
ORIGINNODE	System_Name
TIMESTAMP	Timestamp
USER_NAME	User_Name

UNIXFILE historical table

The UNIXFILE historical table corresponds to the File Information attributes.

Table 49 lists the historical table column heads alphabetically and the corresponding File Information attributes.

Table 49. UNIXFILE table column heads and the corresponding File Information attributes

Historical table column head	Attribute name
ACCESS	Access
ACCESSEDTM	Last_Accessed_Time
CHANGEDTM	Last_Changed_Time
FILE	File
GROUP	Group
LINKNAME	Link_Name
LINKS	Links
ORIGINNODE	System_Name
OWNER	Owner
PATH	Path
SIZE	Size
TIMESTAMP	Timestamp
TYPE	Type

UNIXGROUP historical table

The UNIXGROUP historical table corresponds to the Group attributes.

Table 50 lists the historical table column heads alphabetically and the corresponding Group group attributes.

Table 50. UNIXGROUP table column heads and the corresponding Group attributes

Historical table column head	Attribute name
GRPDUP	Group_Duplicated
GRPID	Group_ID
GRPNAME	Group_Name
ORIGINNODE	System_Name
TIMESTAMP	Timestamp

UNIXIPADDR historical table

The UNIXIPADDR historical table corresponds to the IP Address attributes.

Table 51 lists the historical table column heads alphabetically and the corresponding IP Address group attributes.

Table 51. UNIXIPADDR table column heads and the corresponding IP Address attributes

Historical table column head	Attribute name
DNSNAME	DNS_Name
INTFNAME	Network_Interface_Name
IPADDRESS	IP_Address
IPVERSION	IP_Version

Table 51. UNIXIPADDR table column heads and the corresponding IP Address attributes (continued)

Historical table column head	Attribute name
ORIGINNODE	System_Name
TIMESTAMP	Timestamp

UNIXLPAR historical table

The UNIXLPAR historical table corresponds to the AIX LPAR attributes.

Table 52 lists the historical table column heads alphabetically and the corresponding AIX LPAR group attributes.

Table 52. UNIXLPAR table column heads and the corresponding AIX LPAR attributes

Historical table column head	Attribute name
ACIP	Available_CPUs_in_Pool
ACUIP	Available_CPU_Units_in_Pool
BUSY_PCT	Busy_Pct
CE	CPU_Entitlement
CM	Capped_Mode
CW	Capacity_Weight
DBCP	Donated_Busy_Cycles_Pct
DE	Donation_Enablement
DICP	Donated_Idle_Cycles_Pct
E	Entitlement
EP	Entitlement_Pct
EUP	Entitlement_Used_Pct
HC	Hypervisor_Calls
HOSTNAME	Hostname
LMI	Last_Machine_ID
LN	LPAR_Number
LPAR_NAME	LPAR_Name
MACHINE_ID	Machine_ID
MCCUP	Max_CPU_Cap_Used_Pct
MPC0	Maximum_Pool_Capacity
NOLC	Number_of_Logical_CPUs
NOPC	Number_of_Physical_CPUs
NOPCISP	Number_of_Physical_CPUs_in_Shared_Pool
NOVC	Number_of_Virtual_CPUs
ORIGINNODE	System_Name
PBP	Phys_Busy_Pct
PCSOSP	Physical_CPU_Size_of_Shared_Pool
PCUU	Physical_CPU_Units_Used
PE	Pool_Entitlement

Table 52. UNIXLPAR table column heads and the corresponding AIX LPAR attributes (continued)

Historical table column head	Attribute name
PI	Phantom_Interrupts
POOLID	PoolID
SM	Shared_Mode
SMT_MODE	SMT_Mode
ST	SMT_Threads
TIMESTAMP	Timestamp
TUP	Total_Used_Pct
UCIP	Unallocated_CPU_In_Pool
UPTIME	Uptime
VCCSPS	Virt_Context_CPU_Switches_per_Sec

UNIXMACHIN historical table

The UNIXMACHIN historical table corresponds to the Machine Information attributes.

Table 53 lists the historical table column heads alphabetically and the corresponding Machine Information group attributes.

Table 53. UNIXMACHIN table column heads and the corresponding Machine Information attributes

Historical table column head	Attribute name
HOSTNAME	Hostname
MACSERIAL	Machine_Serial
MODEL	Hardware_Model
ORIGINNODE	System_Name
PHYSPROC	Number_of_Physical_Processors
PMHZ	Processor_Megahertz
TIMESTAMP	Timestamp
UUID	System_Board_UUID
VENDOR	Hardware_Manufacturer
VMID	Virtual_Machine_Identifier

UNIXMEM historical table

The UNIXMEM historical table corresponds to the UNIX Memory attributes.

Table 54 lists the historical table column heads alphabetically and the corresponding UNIX Memory attributes.

Table 54. UNIXMEM table column heads and the corresponding UNIX Memory attributes

Historical table column head	Attribute name
AVAILVM	Avail_Virtual_Storage_MB

Table 54. UNIXMEM table column heads and the corresponding UNIX Memory attributes (continued)

Historical table column head	Attribute name
AVALVMPCT	Virtual_Storage_Pct_Avail
CM	Comp_Memory
DECAY_RATE	Decay_Rate
FREE_PCT	Paging_Space_Free_Pct
MEMAVAIL	Avail_Real_Mem_MB
MEMTOT	Total_Real_Mem_MB
MEMUSED	Used_Real_Mem_MB
NCM	Non_Comp_Memory
ORIGINNODE	System_Name
PRPS	Pages_Read_per_Sec
PSRPS	Paging_Space_Read_per_Sec
PSWPS	Paging_Space_Write_per_Sec
PWPS	Pages_Written_per_Sec
RMAVAP	Avail_Real_Mem_Pct
RMUSDP	Used_Real_Mem_Pct
RR	Repaging_Rate
SWAPAVAIL	Avail_Swap_Space_MB
SWAPTOT	Total_Swap_Space_MB
SWAPUSED	Used_Swap_Space_MB
SWAVAP	Avail_Swap_Space_Pct
SWUSDP	Used_Swap_Space_Pct
TIMESTAMP	Timestamp
USED_PCT	Paging_Space_Used_Pct
USEDVM	Used_Virtual_Storage_MB
USEDVMPCT	Virtual_Storage_Pct_Used
VMPGFAULT	Page_Faults
VMPGIN	Page_Ins
VMPGIN1	Page_In_1Min
VMPGIN5	Page_In_5Min
VMPGIN15	Page_In_15Min
VMPGIN60	Page_In_60Min
VMPGINKBS	Page_In_KB_S
VMPGINREQ	Page_In_Reqs
VMPGOUT	Page_Outs
VMPGOUT1	Page_Out_1Min
VMPGOUT5	Page_Out_5Min
VMPGOUT15	Page_Out_15Min
VMPGOUT60	Page_Out_60Min
VMPGOUTKBS	Page_Out_KB_S

Table 54. UNIXMEM table column heads and the corresponding UNIX Memory attributes (continued)

Historical table column head	Attribute name
VMPGOUTREQ	Page_Out_Reqs
VMPGRCLM	Page_Reclaims
VMSCAN	Page_Scan
VMSCAN1	Page_Scan_1Min
VMSCAN5	Page_Scan_5Min
VMSCAN15	Page_Scan_15Min
VMSCAN60	Page_Scan_60Min
VMSCANKB	Page_Scan_KB
VMTOT	Total_Virtual_Storage_MB

UNIXNET historical table

The UNIXNET historical table corresponds to the Network attributes.

Table 55 lists the historical table column heads alphabetically and the corresponding Network attributes.

Table 55. UNIXNET table column heads and the corresponding Network attributes

Historical table column head	Attribute name
BUP	Bandwidth_Util_Pct
FCOLLSNS	Collisions
FDNSNAME	Interface_DNS_Name
FIBYTES	Received_Count
FIERRORS	Input_Errors
FIFRAMES	Frames_Received
FIPADDR	Interface_IP_Address
FMTU	Transmission_Unit_Maximum
FNAME	Network_Interface_Name
FOBYTES	Transmitted_Count
FOERRORS	Output_Errors
FOFRAMES	Frames_Transmitted
FSTATUS	Interface_Status
FUNIT	Subunit_Driver
ORIGINNODE	System_Name
TIMESTAMP	Timestamp

Note: Frames_Received and Frames_Transmitted refer to packets.

UNIXNFS historical table

The UNIXNFS historical table corresponds to the NFS and RPC Statistics attributes.

Table 56 lists the historical table column heads alphabetically and the corresponding NFS and RPC Statistics attributes.

Table 56. UNIXNSF table column heads and the corresponding NFS and RPC Statistics attributes

Historical table column head	Attribute name
NCBAD	NFS_Client_Calls_Rejected
NCCALLS	NFS_Client_Calls
NCCREATE	NFS_Client_File_Creates
NCFSSTAT	NFS_Client_File_System_Statistics_Calls
NCGETATT	NFS_Client_Get_Attribute_Calls
NCLINK	NFS_Client_Link_Calls
NCLOOKUP	NFS_Client_Lookups
NCMKDIR	NFS_Client_Make_Directory_Calls
NCNULL	NFS_Client_Null_Calls
NCPERC	NFS_Client_Rejected_Calls_Percentage
NCRDDIR	NFS_Client_Read_Directory_Calls
NCRDLINK	NFS_Client_Read_Link_Calls
NCREAD	NFS_Client_Read_Calls
NCREMOVE	NFS_Client_Remove_File_Calls
NCRENAME	NFS_Client_Rename_File_Calls
NCRMDIR	NFS_Client_Remove_Directory_Calls
NCROOT	NFS_Client_root_Calls
NCSETATT	NFS_Client_Set_Attribute_Calls
NCSYMLNK	NFS_Client_Symbolic_Link_Calls
NCWRCACH	NFS_Client_Write_Cache_Calls
NCWRITE	NFS_Client_Writes
NSBAD	NFS_Server_Calls_Rejected
NSCALLS	NFS_Server_Calls
NSCREATE	NFS_Server_File_Creates
NSFSSTAT	NFS_Server_File_System_Statistics_Calls
NSGETATT	NFS_Server_Get_Attribute_Calls
NSLINK	NFS_Server_Link_Calls
NSLOOKUP	NFS_Server_Lookups
NSMKDIR	NFS_Server_Make_Directory_Calls
NSNULL	NFS_Server_Null_Calls
NSRDDIR	NFS_Server_Read_Directory_Calls
NSRDLINK	NFS_Server_Read_Link_Calls
NSREAD	NFS_Server_Read_Calls
NSPERC	NFS_Server_Rejected_Calls_Percentage
NSREMOVE	NFS_Server_Remove_File_Calls
NSRENAME	NFS_Server_Rename_File_Calls
NSRMDIR	NFS_Server_Remove_Directory_Calls

Table 56. UNIXNSF table column heads and the corresponding NFS and RPC Statistics attributes (continued)

Historical table column head	Attribute name
NSROOT	NFS_Server_root_Calls
NSSETATT	NFS_Server_Set_Attribute_Calls
NSSYMLNK	NFS_Server_Symbolic_Link_Calls
NSWRCACH	NFS_Server_Write_Cache_Calls
NSWRITE	NFS_Server_Writes
ORIGINNODE	System_Name
RCAREF	RPC_Client_Times_Authentication_Refreshed
RCBAD	RPC_Client_Calls_Rejected_by_Server
RCBADXID	RPC_Client_Replies_Not_Matching_Calls
RCRETRAN	RPC_Client_Calls_Retransmitted
RCTIMOUT	RPC_Client_Calls_Timed_Out
RCWAIT	RPC_Client_Times_Call_Wait_On_Busy
RSBAD	RPC_Server_Calls_Rejected
RSBADHDR	RPC_Server_Packets_with_Malformed_Header
RSBADLEN	RPC_Server_Packets_Too_Short
RSNULL	RPC_Server_Times_RPC_Packet_Unavailable
TIMESTAMP	Timestamp
ZTITLE	Attribute_Title *Column seen in historical data collection tables but currently not collecting validated data.*
ZVALUE	Attribute_Value *Column seen in historical data collection tables but currently not collecting validated data.*

UNIXOS historical table

The UNIXOS historical table corresponds to the System attributes.

Table 57 lists the historical table column heads alphabetically and the corresponding System attributes.

Table 57. UNIXOS table column heads and the corresponding System attributes

Historical table column head	Attribute name
BOOTTIME	Boot_Time
BREAD	Block_Reads
BWRITE	Block_Writes
CPUBUSY	CPU_Busy
DEVINT	Device_Interrupts
LREAD	Logical_Block_Reads
LWRITE	Logical_Block_Writes

Table 57. UNIXOS table column heads and the corresponding System attributes (continued)

Historical table column head	Attribute name
MEMFREE	DUMMY-Memory_Free *Column seen in historical data collection tables but currently not collecting validated data.*
MEMUSED	Memory_Used *Column seen in historical data collection tables but currently not collecting validated data.*
MDMINT	Modem_Interrupts *Column seen in historical data collection tables but currently not collecting validated data.*
NETADDR	Net_Address
NETCONNECT	Active_Internet_Connections *Column seen in historical data collection tables but currently not collecting validated data.*
NETLOAD1	Load_Average_1_Min
NETLOAD2	Load_Average_5_Min
NETLOAD3	Load_Average_15_Min
NETSOCKET	Active_Sockets *Column seen in historical data collection tables but currently not collecting validated data.*
NOC	Number_of_CPUs
NOSYSPROCS	System_Procs_Number
NOUSRSESS	Users_Session_Number
NSYSTHRD	System_Threads *Column seen in historical data collection tables but currently not collecting validated data.*
ORIGINNODE	System_Name
PC	Physical_Consumption
PENDIOWT	Pending_IO_Waits *Column seen in historical data collection tables but currently not collecting validated data.*
PHREAD	NonBlock_Reads
PHWRITE	NonBlock_Writes
PIDLE	Processes_Idle
PRUNABLE	Processes_Runnable
PRUNNING	Processes_Running
PSLEEPING	Processes_Sleeping
PSTOPPED	Processes_Stopped
PSWITCH	CPU_Context_Switches
PZOMBIE	Processes_Zombie

Table 57. UNIXOS table column heads and the corresponding System attributes (continued)

Historical table column head	Attribute name
RCVINT	Receive_Interrupts *Column seen in historical data collection tables but currently not collecting validated data.*
SBCP	Stolen_Busy_Cycles_Pct
SICP	Stolen_Idle_Cycles_Pct
SSV	System_Software_Version
STARTIO	Start_IO *Column seen in historical data collection tables but currently not collecting validated data.*
SYSCALL	System_Calls
SYSEXEC	Execs_Executed
SYSFORK	Forks_Executed
SYSREAD	System_Read
SYSWRITE	System_Write
SYSTEMTYPE	Type
SYSUPTIME	Up_Time
SYSTEMVERS	Version
THRDRUNQ	Threads_in_Run_Queue *Column seen in historical data collection tables but currently not collecting validated data.*
THRDWAIT	Threads_Waiting *Column seen in historical data collection tables but currently not collecting validated data.*
TIMESTAMP	Timestamp
TOTREALMEM	Total_Real_Memory
TOTVIRTMEM	Total_Virtual_Memory
TSIHP	Time_Spent_in_Hypervisor_Pct
UNIXIDLCPU	Idle_CPU
UNIXSYSCPU	System_CPU
UNIXUSRCPU	User_CPU
UNIXWAITIO	Wait_I/O
UPTIME	UpTime
VMFREEMEM	Free_Memory
VMFREEPRC	Virtual_Memory_Percent_Available
VMFREESWAP	Active_Virtual_Memory
VMINPGWAIT	Processes_Waiting
VMINRUNQ	Processes_in_Run_Queue
VMPGFAULTS	Page_Faults
VMPGIN	Pages_Paged_In
VMPGOUT	Pages_Paged_Out

Table 57. UNIXOS table column heads and the corresponding System attributes (continued)

Historical table column head	Attribute name
VMPGRCLM	Page_Reclaims
VMPGSIN	Page_Ins
VMPGSOUT	Page_Outs
VMSCAN	Page_Scan_Rate
VMUSEDPRC	Virtual_Memory_Percent_Used
XMTINT	Transmit_Interrupts *Column seen in historical data collection tables but currently not collecting validated data.*
ZATTRIB	Parameter *Column seen in historical data collection tables but currently not collecting validated data.*
ZVALUE	Value *Column seen in historical data collection tables but currently not collecting validated data.*

UNIXPING historical table

The UNIXPING historical table corresponds to the Ping attributes.

Table 58 lists the historical table column heads alphabetically and the corresponding Ping group attributes.

Table 58. UNIXPING table column heads and the corresponding Ping attributes

Historical table column head	Attribute name
HOSTRESP	Response_Time
ORIGINNODE	System_Name
PINGHOST	Target_Host
PINGRESULT	Ping_Result
TIMESTAMP	Timestamp

UNIXPRINTQ historical table

The UNIXPRINTQ historical table corresponds to the Print Queue attributes.

Table 59 lists the historical table column heads alphabetically and the corresponding Print Queue group attributes.

Table 59. UNIXPRINTQ table column heads and the corresponding Print Queue attributes

Historical table column head	Attribute name
DEVICENM	Device_Name
ORIGINNODE	System_Name
PRINTQSIZE	Print_Queue_Job_Size
PRNTQDEPTH	Print_Queue_Depth
PRNTQSTATE	Print_Queue_Status

Table 59. UNIXPRINTQ table column heads and the corresponding Print Queue attributes (continued)

Historical table column head	Attribute name
PRT	Print_Queue_Name
TIMESTAMP	Timestamp

UNIXPS historical table

The UNIXPS historical table corresponds with the Process attributes.

Table 60 lists the historical table column heads alphabetically and the corresponding Process attributes.

Table 60. UNIXPS table column heads and the corresponding Process attributes

Historical table column head	Attribute name
ADDR	Entry_Address
CHILDTIME	Total_Child_CPU_Time
CHILDSTIME	Child_System_CPU_Time
CHILDUTIME	Child_User_CPU_Time
CMD	Command
COMMAND	Process_Command
CONTSWITCH	Context_Switch
CPU	CPU_Utilization
CPUID	CPU_ID
CPUPERCENT	CPU_Pct
CPUTIME	CPU_Time *Column seen in historical data collection tables but currently not collecting validated data.*
EGID	Effective_Group_ID
ELAPTIME	Elapsed_Time
EUID	Effective_User_ID
EVENT	Event_Waited_On
EXECSTATE	Execution_State
FLAG	Flag
GID	Real_Group_ID
HEAP	Heap_Size
INVCONTSWT	Involuntary_Context_Switches
MAJORFAULT	Major_Fault
MEMPERCENT	Mem_Pct
MINORFAULT	Minor_Fault
NICE	Nice_Value
ORIGINNODE	System_Name
PID	Process_ID
PGID	Process_Group_Leader_ID

Table 60. UNIXPS table column heads and the corresponding Process attributes (continued)

Historical table column head	Attribute name
PPID	Parent_Process_ID
PRIORITY	Priority
PSU	Page_Space_Used
RDS	Resident_Data_Size
RTS	Resident_Text_Size
READWRITE	Read/Write
REPTYPE	Reptype *Column seen in historical data collection tables. IBM internal use only.*
SAMPCPUTCT	Sample_CPU_Pct *Column seen in historical data collection tables but currently not collecting validated data.*
SCHEDCLASS	Scheduling_Class
SESSIONID	Session_ID
SIZE	Size
STACK	Stack_Size
STARTTIME	StartTime
SYSTEMTIM	System_CPU_Time
THREADCNT	Thread_Count
TIME	Time
TIMESTAMP	Timestamp
TOTCPUPERC	Total_CPU_Percent
TOTALTIME	Total_CPU_Time
TTY	Terminal_Device
UID	User_ID
USERNAME	User_Name
USERTIME	User_CPU_Time
VSIZE	Virtual_Size
WAITCPUTIM	Wait_CPU_Time
WAITLTIME	Wait_Lock_Time
WLM_NAME	WLM_Name
WPAR_NAME	WPAR_Name
ZATTRIB	Parameter *Column seen in historical data collection tables but currently not collecting validated data.*
ZVALUE	Value *Column seen in historical data collection tables but currently not collecting validated data.*

UNIXSOLZON historical table

The UNIXSOLZON historical table corresponds to the Solaris Zones attributes.

Table 61 lists the historical table column heads alphabetically and the corresponding Solaris Zones group attributes.

Table 61. UNIXSOLZON table column heads and the corresponding Solaris Zones attributes

Historical table column head	Attribute name
CAPCPU	Capped_CPU
CAPMEM	Capped_Memory
CPUSHARES	CPU_Shares
DEDCPU	Dedicated_CPU
IPID	Init_PID
ORIGINNODE	System_Name
POOLID	Pool_ID
SCHED	Scheduler
SHAREPCT	CPU_Share_PCT
TIMESTAMP	Timestamp
ZCPU	Zone_CPU_Usage
ZCPUS	Total_CPUs
ZID	Zone_ID
ZONENAME	Name
ZPATH	Path
ZRSS	Physical_Memory
ZSTATUS	Status
ZVMS	Virtual_Memory

UNIXTCP historical table

The UNIXTCP historical table corresponds to the TCP Statistics attributes.

Table 62 lists the historical table column heads alphabetically and the corresponding TCP Statistics group attributes.

Table 62. UNIXTCP table column heads and the corresponding TCP Statistics attributes

Historical table column head	Attribute name
ORIGINNODE	System_Name
PKTRETRPS	TCP_Data_Packets_Retrans_per_sec
TIMESTAMP	Timestamp

UNIXUSER historical table

The UNIXUSER historical table corresponds to the User group attributes.

Table 63 lists the historical table column heads alphabetically and the corresponding User attributes.

Table 63. UNIXUSER table column heads and the corresponding User attributes

Historical table column head	Attribute name
ORIGINNODE	System_Name
PID	Process_ID *Column seen in historical data collection tables but currently not collecting validated data.*
TIMESTAMP	Timestamp
UID	User_ID
USERIDLE	Idle_Time
USERLOGIN	Login_Name
USERNAME	Name
USERSITE	Location
USERTTY	Terminal
USERWHEN	Login_Time

UNIXWPARCP historical table

The UNIXWPARCP historical table corresponds to the AIX WPAR CPU attributes.

Table 64 lists the historical table column heads alphabetically and the corresponding AIX WPAR CPU group attributes.

Table 64. UNIXWPARCP table column heads and the corresponding AIX WPAR CPU attributes

Historical table column head	Attribute name
CCL	CPU_Consumption_Limit
LCCP	LPAR_CPU_Consumed_Pct
LE	LPAR_Entitlement
NCC	Num_CPUs_Consumed
ORIGINNODE	System_Name
RCLHM	RC_CPU_Limits_Hard_Max
SCP	System_CPU_Pct
TIMESTAMP	Timestamp
UCP	User_CPU_Pct
WCCP	WPAR_CPU_Consumed_Pct
WPAR_NAME	WPAR_Name

UNIXWPARFS historical table

The UNIXWPARFS historical table corresponds to the AIX WPAR File System attributes.

Table 65 on page 293 lists the historical table column heads alphabetically and the corresponding AIX WPAR File System group attributes.

Table 65. UNIXWPARFS table column heads and the corresponding AIX WPAR File System attributes

Historical table column head	Attribute name
DN	Device_Name
MO	Mount_Options
MP	Mount_Point
NODE_NAME	Node_Name
ORIGINNODE	System_Name
TIMESTAMP	Timestamp
VFS_TYPE	VFS_Type
WPAR_NAME	WPAR_Name

UNIXWPARIN historical table

The UNIXWPARIN historical table corresponds to the AIX WPAR attributes.

Table 66 lists the historical table column heads alphabetically and the corresponding AIX WPAR group attributes.

Table 66. UNIXWPARIN table column heads and the corresponding AIX WPAR Information attributes

Historical table column head	Attribute name
AO	Admin_Operation
API	Admin_Process_ID
AST	Admin_Start_Time
AUTOSTART	Autostart
C	Checkpointable
C_WHSC	C_WHSC
HOME	Home
HOSTNAME	Hostname
IP_ADDRESS	IP_Address
ORIGINNODE	System_Name
OWNER	Owner
RC_RSET	RC_RSet
RCLHM	RC_CPU_Limits_Hard_Max
RCLM	RC_CPU_Limits_Min
RCLSM	RC_CPU_Limits_Soft_Max
RCS	RC_CPU_Shares
RIA	RC_Is_Active
RMLHM	RC_Memory_Limits_Hard_Max
RMLM	RC_Memory_Limits_Min
RMLSM	RC_Memory_Limits_Soft_Max
RMP	RC_Max_Processes
RMS	RC_Memory_Shares

Table 66. UNIXWPARIN table column heads and the corresponding AIX WPAR Information attributes (continued)

Historical table column head	Attribute name
RMT	RC_Max_Threads
RPPVL	RC_per_Process_VM_Limit
STATE	State
SUD	Shares_usr_Dir
TIMESTAMP	Timestamp
TYPE	Type
WPAR_NAME	WPAR_Name
WAP	WPAR_Application_Path

UNIXWPARNE historical table

The UNIXWPARNE historical table corresponds to the AIX WPAR Network attributes.

Table 67 lists the historical table column heads alphabetically and the corresponding AIX WPAR Network group attributes.

Table 67. UNIXWPARNE table column heads and the corresponding AIX WPAR Network attributes

Historical table column head	Attribute name
BI	Broadcast_IP
IN	Interface_Name
IP_ADDRESS	IP_Address
NM	Network_Mask
ORIGINNODE	System_Name
TIMESTAMP	Timestamp
WPAR_NAME	WPAR_Name

UNIXWPARPM historical table

The UNIXWPARPM historical table corresponds to the AIX WPAR Physical Memory attributes.

Table 68 lists the historical table column heads alphabetically and the corresponding AIX WPAR Physical Memory group attributes.

Table 68. UNIXWPARPM table column heads and the corresponding AIX WPAR Physical Memory attributes

Historical table column head	Attribute name
FMM	Free_Memory_MB
FMP	Free_Memory_Pct
MSM	Memory_Size_MB
LMSM	LPAR_Memory_Size_MB
LMUP	LPAR_Memory_Used_Pct

Table 68. UNIXWPARPM table column heads and the corresponding AIX WPAR Physical Memory attributes (continued)

Historical table column head	Attribute name
ORIGINNODE	System_Name
RMLHM	RC_Memory_Limits_Hard_Max
TIMESTAMP	Timestamp
UMM	Used_Memory_MB
UMP	Used_Memory_Pct
WPAR_NAME	WPAR_Name

Appendix E. Monitoring Agent for UNIX data collection

In general, the Monitoring Agent for UNIX gathers data when requested to satisfy a workspace refresh, situation sampling of attributes, or historical data collection. All attributes in the attribute groups that make up a workspace or situation are gathered at that time. The default refresh/sampling intervals were chosen such that the agent does not put a significant load on the system as it gathers the data. The following tables list the mechanisms used to gather attribute values for each attribute group:

Table 69. Mechanisms used to gather FILEINFO attributes

Attribute	Collection methods	AIX	HPUX	Solaris
ACCESS	API	stat()	stat()	stat()
ACCESSEDTM	API	stat()	stat()	stat()
ACCESSTM	API	stat()	stat()	stat()
CHANGEDTM	API	stat()	stat()	stat()
CHECKSUM	API	ICC library	ICC library	ICC library
FILE	API	From Situation/Query	From Situation/Query	From Situation/Query
GROUP	API	stat()	stat()	stat()
HASHALGO	API	From Situation/Query	From Situation/Query	From Situation/Query
LINKNAME	API	stat()	stat()	stat()
LINKS	API	stat()	stat()	stat()
MODE	API	stat()	stat()	stat()
MODRESULT	API	ICC library	ICC library	ICC library
OWNER	API	stat()	stat()	stat()
PATH	API	From Situation/Query | "/"	From Situation/Query | "/"	From Situation/Query | "/"
SIZE	API	stat()	stat()	stat()
SIZEMB	API	stat()	stat()	stat()
TYPE	API	stat()	stat()	stat()
UFILE	API	stat()	stat()	stat()
UGROUP	API	stat()	stat()	stat()
ULINKNAME	API	stat()	stat()	stat()
UOWNER	API	stat()	stat()	stat()
UPATH	API	stat()	stat()	stat()

Table 70. Mechanisms used to gather UNIXALLUSR attributes

Attribute	Collection methods	AIX	HPUX	Solaris
ORIGINNODE	API	getpwent	getpwent	getpwent
PWNULL	API	getpwent	getpwent	getpwent
TIMESTAMP	API	getpwent	getpwent	getpwent

Table 70. Mechanisms used to gather UNIXALLUSR attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
UID	API	getpwent	getpwent	getpwent
USERDUP	API	getpwent	getpwent	getpwent
USERNAME	API	getpwent	getpwent	getpwent
USERSES	API	getutxent	getutxent	getutxent

Table 71. Mechanisms used to gather UNIXAMS attributes

Attribute	Collection methods	AIX	HPUX	Solaris
AME	SPMI	LPAR/iome	N/A	N/A
AMEI	SPMI	LPAR/iomu	N/A	N/A
AML	SPMI	LPAR/meml	N/A	N/A
AMS_MODE	SPMI	LPAR/vrmenabled	N/A	N/A
API	SPMI	LPAR/mpid	N/A	N/A
APM	SPMI	LPAR/pmem	N/A	N/A
APS	SPMI	LPAR/mpsz	N/A	N/A
HPI	SPMI	LPAR/hpi	N/A	N/A
HPIT	SPMI	LPAR/hpit	N/A	N/A

Table 72. Mechanisms used to gather UNIXCPU attributes

Attribute	Collection methods	AIX	HPUX	Solaris
AVCPUBIZ1	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
AVCPUBIZ5	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
AVCPUBIZ15	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
AVCPUBIZ60	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
AVCPUSYS1	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
AVCPUSYS5	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
AVCPUSYS15	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
AVCPUSYS60	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
AVCPUUSR1	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
AVCPUUSR5	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
AVCPUUSR15	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat

Table 72. Mechanisms used to gather UNIXCPU attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
AVCPUUSR60	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
CPUBUSY	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
CPUID	API	5.2: perfstast_cpu 5.3: mpstat	mpctl	kstat
CPUSTAT	API	5.2: perfstast_cpu 5.3: mpstat	get_ sysinfo	kstat
CPUUSAGE	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
CPUUVS	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
CSPS	API	ptx_get_cpuinfo (delta of cswitches / elapsed time)	N/A	N/A
CSW	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
ICSW	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
IDLECPU	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
INTRRUPT	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
INTRTHRD	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
LCS	API	ptx_get_cpuinfo (delta of cswitches / elapsed time)	N/A	N/A
MAJF	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
MINF	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
PC	API	ptx_get_cpuinfo ((usr + kernel + wait + idle) / tb_last)	N/A	N/A
SMTX	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
SRWLOCKS	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
SYSCALL	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
SYSCPU	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
THRDMIGR	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat
USRCPU	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat

Table 72. Mechanisms used to gather UNIXCPU attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
WAITIO	API	5.2: perfstast_cpu 5.3: mpstat	get_sysinfo	kstat
XCALLS	API	5.2: perfstast_cpu 5.3: mpstat	pstat_ getvminfo	kstat

Table 73. Mechanisms used to gather UNIXDEVIC attributes

Attribute	Collection methods	AIX	HPUX	Solaris
DCLASS	AIX Script Data Provider	/usr/sbin/lsdev -C -F	N/A	N/A
NAME	AIX Script Data Provider	/usr/sbin/lsdev -C -F	N/A	N/A
PARENT	AIX Script Data Provider	/usr/sbin/lsdev -C -F	N/A	N/A
STATE	AIX Script Data Provider	/usr/sbin/lsdev -C -F	N/A	N/A
TYPE	AIX Script Data Provider	/usr/sbin/lsdev -C -F	N/A	N/A

Table 74. Mechanisms used to gather UNIXDISK attributes

Attribute	Collection methods	AIX	HPUX	Solaris
DSKNAME	API	statvfs	statvfs	statvfs64
DSKSIZE	API	statvfs	statvfs	statvfs64
DSKSIZEGB	API	statvfs	statvfs	statvfs64
DSKSIZEMB	API	statvfs	statvfs	statvfs64
FILESYSTYP	API	statvfs64()	statvfs64()	statvfs64()
FSSTATUS	API	UP or DOWN (if statfs64() times out)	UP or DOWN (if statfs64() times out)	UP or DOWN (if statfs64() times out)
FSTYPE	API	statvfs	statvfs	statvfs64
INODEFREE	API	statvfs	statvfs	statvfs64
INODESIZE	API	statvfs	statvfs	statvfs64
INODEUSED	API	statvfs	statvfs	statvfs64
MOUNTPT	API	statvfs	statvfs	statvfs64
PCTINDAVAL	API	statvfs	statvfs	statvfs64
PCTINDUSED	API	statvfs	statvfs	statvfs64
PCTSPCAV	API	statvfs	statvfs	statvfs64
PCTSPCUSED	API	statvfs	statvfs	statvfs64
SPAVGB	API	statvfs	statvfs	statvfs64

Table 74. Mechanisms used to gather UNIXDISK attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
SPAVMB	API	statvfs	statvfs	statvfs64
SPCAVAIL	API	statvfs	statvfs	statvfs64
SPCUSED	API	statvfs	statvfs	statvfs64
SPUSEDGB	API	statvfs	statvfs	statvfs64
SPUSEDMB	API	statvfs	statvfs	statvfs64
UDSKNAME	API	statvfs	statvfs	statvfs64
UMOUNTPT	API	statvfs	statvfs	statvfs64
VGN	AIX Script Data Provider	/usr/sbin/lsvg -l	N/A	N/A

Table 75. Mechanisms used to gather UNIXDPERF attributes

Attribute	Collection methods	AIX	HPUX	Solaris
ASS	API	delta of dk_q_sampled / CPU_Time	N/A	N/A
AVGSERV	API	/bin/iostat	pstat_getdisk	kstat
DSKAVQUE	API	/bin/iostat	pstat_getdisk	kstat
DSKAVWAIT	API	/bin/iostat	pstat_getdisk	kstat
DSKAVXFR	API	/bin/iostat	pstat_getdisk	kstat
DSKBUSY	API	/bin/iostat	pstat_getdisk	kstat
DSKBYTESIZ	API	/bin/iostat	pstat_getdisk	kstat
DSKNAME	API	/bin/iostat	pstat_getdisk	kstat
DSKRDBYSEC	API	/bin/iostat	pstat_getdisk	kstat
DSKRDSEC	API	/bin/iostat	pstat_getdisk	kstat
DSKWRBYSEC	API	/bin/iostat	pstat_getdisk	kstat
DSKWRSEC	API	/bin/iostat	pstat_getdisk	kstat
DSKXFERRAT	API	/bin/iostat	pstat_getdisk	kstat
MOUNTPT	API	/bin/iostat	pstat_getdisk	kstat
PARENT	API	“N/A”	N/A	N/A
PCTDSKRD	API	/bin/iostat	pstat_getdisk	kstat
PCTDSKWR	API	/bin/iostat	pstat_getdisk	kstat
SFPS	API	delta of dk_q_full / Elapsed_Time	N/A	N/A
TKPS	API	delta of dk_xfers / (1024.0 * Elapsed_Time)	N/A	N/A
TYPE	API	"Adapter" | "Disk" | “Path/Disk" | "Vtarget/Disk" | etc	N/A	N/A
UDSKNAME	API	/bin/iostat	pstat_getdisk	kstat

Table 76. Mechanisms used to gather UNIXDUSERS attributes

Attribute	Collection methods	AIX	HPUX	Solaris
AL	AIX Script Data Provider	lsuser -c ALL	N/A	N/A
EXPIRES	AIX Script Data Provider	lsuser -c ALL	N/A	N/A
L	AIX Script Data Provider	lsuser -c ALL	N/A	N/A
ROLES	AIX Script Data Provider	lsuser -c ALL	N/A	N/A
USERNAME	AIX Script Data Provider	lsuser -c ALL	N/A	N/A

Table 77. Mechanisms used to gather UNIXFILCMP attributes

Attribute	Collection methods	AIX	HPUX	Solaris
COMPOPT	API	From Situation/Query	From Situation/Query	From Situation/Query
COMPRESULT	API	/usr/bin/diff | /usr/bin/cmp	/usr/bin/diff | /usr/bin/cmp	/usr/bin/diff | /usr/bin/cmp
FILENAME1	API	From Situation/Query	From Situation/Query	From Situation/Query
FILENAME2	API	From Situation/Query	From Situation/Query	From Situation/Query

Table 78. Mechanisms used to gather UNIXFILPAT attributes

Attribute	Collection methods	AIX	HPUX	Solaris
FILENAME	API	From Situation/Query	From Situation/Query	From Situation/Query
MATCHCNT	API	/usr/bin/grep	/usr/bin/grep	/usr/bin/grep
MATCHOPT	API	From Situation/Query	From Situation/Query	From Situation/Query
MATCHPAT	API	From Situation/Query	From Situation/Query	From Situation/Query

Table 79. Mechanisms used to gather UNIXGROUP attributes

Attribute	Collection methods	AIX	HPUX	Solaris
GRPDUP	API	getgrent	getgrent	getgrent
GRPID	API	getgrent	getgrent	getgrent
GRPNAME	API	getgrent	getgrent	getgrent

Table 80. Mechanisms used to gather UNIXIPADDR attributes

Attribute	Collection methods	AIX	HPUX	Solaris
DNSNAME	API	getaddrinfo	getaddrinfo	getaddrinfo

Table 80. Mechanisms used to gather UNIXIPADDR attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
INTFNAME	API	SIOCGIFCONF	SIOCGIF CONF/ SIOCGLIF CONF	SIOCGLIF CONF
IPADDRESS	API	SIOCGIFCONF	SIOCGIF CONF/ SIOCGLIF CONF	SIOCGLIF CONF
IP_Version	API	socket (SOCK_DGRAM)	socket (SOCK_DGRAM)	socket (SOCK_DGRAM)

Table 81. Mechanisms used to gather UNIXLPAR attributes

Attribute	Collection methods	AIX	HPUX	Solaris
ACIP	SPMI	LPAR/app	N/A	N/A
ACUIP	SPMI	LPAR/app * 100	N/A	N/A
BUSY_PCT	SPMI	LPAR/lbusy	N/A	N/A
CE	SPMI	LPAR/entitledcap	N/A	N/A
CM	SPMI	LPAR/capped	N/A	N/A
CW	SPMI	LPAR/varwght	N/A	N/A
DBCP	SPMI	LPAR/%bdon	N/A	N/A
DE	SPMI	LPAR/donate_enabled	N/A	N/A
DICP	SPMI	LPAR/%idon	N/A	N/A
E	SPMI	If dedicated, then LPAR/vcpu; otherwise, LPAR/ent.	N/A	N/A
EP	SPMI	LPAR/entpct	N/A	N/A
EUP	SPMI	If dedicated, then (LPAR/physc / LPAR/vcpu * 100); otherwise, LPAR/entc.	N/A	N/A
HC	SPMI	LPAR/hcalls	N/A	N/A
LPAR_NAME	API	sysconfig	N/A	N/A
HOSTNAME	SPMI	uname	N/A	N/A
LMI	SPMI	previous value of unamex	N/A	N/A
LN	SPMI	LPAR/lparnum	N/A	N/A
MACHINE_ID	SPMI	unamex	N/A	N/A
MCCUP	SPMI	If dedicated or uncapped, then (LPAR/pbusy / LPAR/vcpu); otherwise, (LPAR/pbusy / LPAR/ent).	N/A	N/A

Table 81. Mechanisms used to gather UNIXLPAR attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
MPC0	SPMI	LPAR/max_pool_cap	N/A	N/A
NOLC	SPMI	LPAR/lcpu	N/A	N/A
NOPC	SPMI	LPAR/pcpu	N/A	N/A
NOPCISP	SPMI	LPAR/pcpuinpool	N/A	N/A
NOVC	SPMI	LPAR/vcpu	N/A	N/A
PBP	SPMI	LPAR/pbusy	N/A	N/A
PCSOSP	SPMI	LPAR/pcpuinpool * 100	N/A	N/A
PCUU	SPMI	LPAR/physc * 100	N/A	N/A
PE	SPMI	LPAR/pool_entc	N/A	N/A
PI	SPMI	LPAR/phint	N/A	N/A
POOLID	SPMI	LPAR/poolid	N/A	N/A
SM	SPMI	LPAR/shared	N/A	N/A
SMT_MODE	SPMI	LPAR/smt	N/A	N/A
ST	SPMI	LPAR/smtctl	N/A	N/A
TUP	SPMI	LPAR/ent * LPAR/entc / LPAR/pcpu	N/A	N/A
UCIP	SPMI	LPAR/unalloccap	N/A	N/A
UPTIME	SPMI	time since boot	N/A	N/A
VCCSPS	SPMI	LPAR/vcsw	N/A	N/A

Table 82. Mechanisms used to gather UNIXMACHIN attributes

Attribute	Collection methods	AIX	HPUX	Solaris
HOSTNAME	API	gethostname()	gethostname()	gethostname()
MACSERIAL	API	uname -uM	confstr (_CS_MACHINE _SERIAL)	/usr/sbin/eeprom nvramrc | /opt/SUNWsneep/ bin/sneep*
MODEL	API	getaddr()	_CS_MACHINE_ MODEL	SI_PLATFORM
PHYSPROC	API	perfstat_cpu	pstat_ getdynamic	_SC_ NPROCESSORS _CONF
PMHZ	API	perfstat_cpu _total	_SC_CLK_TCK	kstat()
UUID	API	uname -f	confstr (_CS_MACHINE _IDENT)	gethostid()
VENDOR	API	set to "IBM"	set to "Hewlett-Packard"	SI_HW_PROVIDER

Table 82. Mechanisms used to gather UNIXMACHIN attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
VMID	API	sysconfig (SYS_GETLPAR _INFO)	confstr (_CS_PARTITION _IDENT)	getzonenamebyid(), getzoneid()

Table 83. Mechanisms used to gather UNIXMEM attributes

Attribute	Collection methods	AIX	HPUX	Solaris
ARCSIZE	API	N/A	N/A	kstat
AVAILVM	API	perfstat_ memory_ total	pstat_get dynamic+ /usr/sbin/swapinfo	swapctl +sysconf
AVALVMPCT	API	perfstat_ memory_ total	pstat_get dynamic+ /usr/sbin/swapinfo	swapctl +sysconf
CM	SPMI	Mem/Real/comp	N/A	N/A
DECAY_RATE	SPMI	Mem/Real/pdecay	N/A	N/A
FREE_PCT	SPMI	PagSp/%totalfree	N/A	N/A
MEMAVAIL	API	perfstat_ memory_ total	pstat_ getdynamic	sysconf
MEMTOT	API	perfstat_ memory_ total	pstat_ getdynamic	sysconf
MEMUSED	API	perfstat_ memory_ total	pstat_ getdynamic	sysconf
NCM	SPMI	Mem/Real/noncomp	N/A	N/A
NETMEMFPCT	API	N/A	N/A	kstat
NETMEMFREE	API	N/A	N/A	kstat
NETMEMUPCT	API	N/A	N/A	kstat
NETMEMUSED	API	N/A	N/A	kstat
PRPS	SPMI	Mem/Virt/pagein	N/A	N/A
PSRPS	SPMI	Mem/Virt/pgspgin	N/A	N/A
PSWPS	SPMI	Mem/Virt/pgspgout	N/A	N/A
PWPS	SPMI	Mem/Virt/pageout	N/A	N/A
RMAVAP	API	perfstat_ memory_ total	pstat_ getdynamic	sysconf
RMUSDP	API	perfstat_ memory_ total	pstat_ getdynamic	sysconf
RR	SPMI	Mem/Real/sysrepag	N/A	N/A
SWAPAVAIL	API	perfstat_ memory_ total	/usr/sbin/swapinfo	swapctl

Table 83. Mechanisms used to gather UNIXMEM attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
SWAPTOT	API	perfstat_ memory_ total	/usr/sbin/swapinfo	swapctl
SWAPUSED	API	perfstat_ memory_ total	/usr/sbin/swapinfo	swapctl
SWAVAP	API	perfstat_ memory_ total	/usr/sbin/swapinfo	swapctl
SWUSDP	API	perfstat_ memory_ total	/usr/sbin/swapinfo	swapctl
USEDVM	API	perfstat_ memory_ total	pstat_get dynamic+ /usr/sbin/swapinfo	swapctl +sysconf
USEDVMPCT	API	perfstat_ memory_ total	pstat_get dynamic+ /usr/sbin/swapinfo	swapctl +sysconf
USED_PCT	SPMI	PagSp/%totalused	N/A	N/A
VMTOT	API	perfstat_ memory_ total	pstat_get dynamic+ /usr/sbin/swapinfo	swapctl +sysconf
VMPGFAULT	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGIN	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGIN1	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGIN5	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGIN15	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGIN60	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGINKBS	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGINREQ	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGOUT	API	perfstat_ memory_ total	pstat_ getvminfo	kstat

Table 83. Mechanisms used to gather UNIXMEM attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
VMPGOUT1	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGOUT5	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGOUT15	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGOUT60	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGOUTKBS	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGOUTREQ	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGRCLM	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMSCAN	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMSCAN1	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMSCAN5	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMSCAN15	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMSCAN60	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMSCANKB	API	perfstat_ memory_ total	pstat_ getvminfo	kstat

Table 84. Mechanisms used to gather UNIXNET attributes

Attribute	Collection methods	AIX	HPUX	Solaris
AVCOL1	API	perfstat_ netinterface	DLPI	kstat
AVCOL5	API	perfstat_ netinterface	DLPI	kstat
AVCOL15	API	perfstat_ netinterface	DLPI	kstat

Table 84. Mechanisms used to gather UNIXNET attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
AVCOL60	API	perfstat_ netinterface	DLPI	kstat
AVGINS1	API	perfstat_ netinterface	DLPI	kstat
AVGINS5	API	perfstat_ netinterface	DLPI	kstat
AVGINS15	API	perfstat_ netinterface	DLPI	kstat
AVGINS60	API	perfstat_ netinterface	DLPI	kstat
AVGINERR1	API	perfstat_ netinterface	DLPI	kstat
AVGINERR5	API	perfstat_ netinterface	DLPI	kstat
AVGINERR15	API	perfstat_ netinterface	DLPI	kstat
AVGINERR60	API	perfstat_ netinterface	DLPI	kstat
AVGOUT1	API	perfstat_ netinterface	DLPI	kstat
AVGOUT5	API	perfstat_ netinterface	DLPI	kstat
AVGOUT15	API	perfstat_ netinterface	DLPI	kstat
AVGOUT60	API	perfstat_ netinterface	DLPI	kstat
BUP	AIX Script Data Provider	/usr/bin/entstat -d (Bytes Sent + Bytes Received / Media Speed Running)	N/A	N/A
FCOLLSNS	API	perfstat_ netinterface	DLPI	kstat
FDNSNAME	API	gethostbyaddr	gethostbyaddr	gethostbyaddr
FDNSNAME4	API	gethostbyaddr	gethostbyaddr	gethostbyaddr
FDSTATUS	API	perfstat_ netinterface	SOCK_ DGRAM ioctl	SOCK_ DGRAM ioctl
FIBYTES	API	perfstat_ netinterface	DLPI	kstat
FIERRORS	API	perfstat_ netinterface	DLPI	kstat
FIFRAMES	API	perfstat_ netinterface	DLPI	kstat
FIPADDR	API	perfstat_ netinterface	SOCK_ DGRAM ioctl	SOCK_ DGRAM ioctl
FMTU	API	perfstat_ netinterface	DLPI	SOCK_ DGRAM ioctl

Table 84. Mechanisms used to gather UNIXNET attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
FNAME	API	perfstat_ netinterface	DLPI	SOCK_ DGRAM ioctl
FOBYTES	API	perfstat_ netinterface	DLPI	kstat
FOERRORS	API	perfstat_ netinterface	DLPI	kstat
FOFRAMES	API	perfstat_ netinterface	DLPI	kstat
FUNIT	API	perfstat_ netinterface	DLPI	SOCK_DGRAM ioctl
IFTYPE	API	perfstat_ netinterface	DLPI	kstat
INMB	API	perfstat_ netinterface	DLPI	kstat
INMBTTL	API	perfstat_ netinterface	DLPI	kstat
MAC_Address	API	/bin/netstat -in	DLPI - DL_PHYS_ ADDR_REQ	/sbin/ifconfig -a
OUTMB	API	perfstat_ netinterface	DLPI	kstat
OUTMBTTL	API	perfstat_ netinterface	DLPI	kstat
PKTCOLPCT	API	perfstat_ netinterface	DLPI	kstat
PKTINERRPT	API	perfstat_ netinterface	DLPI	kstat
PKTOTERRPT	API	perfstat_ netinterface	DLPI	kstat

Table 85. Mechanisms used to gather UNIXNFS attributes

Attribute	Collection methods	AIX	HPUX	Solaris
CRPBADREP	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
CRPBADREPL	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
CRPREJSPCT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
CRPRETLPCT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
CRPRETRPCT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
CRPTIMOUT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCBAD	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCCALLS	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCCREATE	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCFSSTAT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCGETATT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCLINK	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat

Table 85. Mechanisms used to gather UNIXNFS attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
NCLOOKUP	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCMKDIR	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCNULL	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCPERC	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCRDLINK	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCREAD	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCREMOVE	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCRENAME	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCRMDIR	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCROOT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCSETATT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCSYMLNK	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCWRCACH	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NCWRITE	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NFSATRPCT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NFSRDLKPCT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NFSRDPCT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NFSVER	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NFSWRPCT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSBAD	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSCALLS	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSCREATE	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSFSSTAT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSGETATT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSLINK	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSLOOKUP	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSMKDIR	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSNULL	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSPERC	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSRDDIR	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSRDLINK	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSREAD	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSREMOVE	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSRENAME	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSRMDIR	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSROOT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSSETATT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSSYMLNK	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
NSWRCACH	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat

Table 85. Mechanisms used to gather UNIXNFS attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
NSWRITE	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
RCAREF	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
RCBAD	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
RCBADXID	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
RCRETRAN	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
RCTIMOUT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
RCWAIT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
RSBAD	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
RSBADHDR	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
RSBADLEN	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
RSNULL	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
SRPCALLS	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
SRPDUPCHKS	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
SRPDUPREQP	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
SRPREJPCT	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
ZATTRIB	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat
ZVALUE	API	/usr/etc/nfsstat	/usr/bin/nfsstat	kstat

Table 86. Mechanisms used to gather UNIXOS / SP2OS attributes

Attribute	Collection methods	AIX	HPUX	Solaris
AVPGINS1	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
AVPGINS5	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
AVPGINS15	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
AVPGINS60	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
AVPGOUT1	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
AVPGOUT5	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
AVPGOUT15	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
AVPGOUT60	API	perfstat_ memory_ total	pstat_ getvminfo	kstat

Table 86. Mechanisms used to gather UNIXOS / SP2OS attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
AVPGSCAN1	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
AVPGSCAN5	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
AVPGSCAN15	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
AVPGSCAN60	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
AVPRRUNQ60	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
BOOTTIME	API	getutxid	getutxid	getutxid
BREAD	API	perfstat_cpu_ total	get_sysinfo	kstat
BWRITE	API	perfstat_cpu_ total	get_sysinfo	kstat
CPUBUSY	API	perfstat_cpu_ total	get_sysinfo	kstat
DEVINT	API	perfstat_cpu_ total	pstat_ getvminfo	kstat
LREAD	API	perfstat_cpu_ total	get_sysinfo	kstat
LWRITE	API	perfstat_cpu_ total	get_sysinfo	kstat
MDMINT	API	perfstat_cpu_ total	get_sysinfo	kstat
MEMFREE	API	perfstat_ memory_ total	pstat_ getdynamic	sysconf
MEMUSED	API	perfstat_ memory_ total	pstat_ getdynamic	swapctl
NETADDR	API	gethostname() +getaddrinfo	gethostname() +getaddrinfo	gethostname() +getaddrinfo
NETADDR6	API	gethostname() +getaddrinfo	gethostname() +getaddrinfo	gethostname() +getaddrinfo
NETLOAD1	API	perfstat_cpu_ total	pstat_ getdynamic	getloadavg
NETLOAD2	API	perfstat_cpu_ total	pstat_ getdynamic	getloadavg
NETLOAD3	API	perfstat_cpu_ total	pstat_ getdynamic	getloadavg
NOC	SPMI	LPAR/lcpu	N/A	N/A

Table 86. Mechanisms used to gather UNIXOS / SP2OS attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
NOSYSPROCS	API	getprocs64	pstat_getproc	/proc/
NOUSRSESS	API	getutxent	getutxent	getutxent
PC	SPMI	LPAR/physc	N/A	N/A
PGINRATE	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
PGOUTRATE	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
PGSCANRATE	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
PHREAD	API	perfstat_cpu_ total	get_sysinfo	kstat
PHWRITE	API	perfstat_cpu_ total	get_sysinfo	kstat
PIDLE	API	/proc ioctl()	get_sysinfo	pstat_getproc
PRUNABLE	API	getprocs64	get_sysinfo	kstat
PRUNNING	API	N/A	/proc ioctl()	pstat_getproc
PSLEEPING	API	/proc ioctl()	get_sysinfo	pstat_getproc
PSTOPPED	API	/proc ioctl()	get_sysinfo	pstat_getproc
PSWITCH	API	perfstat_cpu_ total	get_sysinfo	kstat
PZOMBIE	API	/proc ioctl()	get_sysinfo	pstat_getproc
RCVINT	API	perfstat_cpu_ total	get_sysinfo	kstat
SBCP	SPMI	LPAR/%bstol	N/A	N/A
SICP	SPMI	LPAR/%istol	N/A	N/A
SSV	SPMI	AIX “/usr/bin/oslevel -s 2>/dev/null”; VIOS, HMC ""/usr/ios/cli/ioscli ioslevel 2>/dev/null"	N/A	N/A
SWAPFREE	API	perfstat_ memory_ total	/usr/sbin/swapinfo	swapctl
SYSCALL	API	perfstat_cpu_ total	get_sysinfo	kstat
SYSEXEC	API	perfstat_cpu_ total	get_sysinfo	kstat
SYSFORK	API	perfstat_cpu_ total	get_sysinfo	kstat
SYSREAD	API	perfstat_cpu_ total	get_sysinfo	kstat
SYSTEMTYPE	API	uname()	uname()	uname()

Table 86. Mechanisms used to gather UNIXOS / SP2OS attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
SYSTEMVERS	API	uname()	uname()	uname()
SYSUPTIME	API	time()-boottime()	time()-boottime()	time()-boottime()
SYSWRITE	API	perfstat_cpu_ total	get_sysinfo	kstat
TOTREALMEM	API	perfstat_ memory_ total	pstat_ getdynamic	sysconf
TOTVIRTMEM	API	perfstat_ memory_ total	pstat_ getdynamic+ /usr/sbin/swapinfo	swapctl+sysconf
TSIHP	SPMI	LPAR/hyppct	N/A	N/A
UNIXIDLCPU	API	5.2 perfstat_cpu 5.3 /usr/bin/mpstat	get_sysinfo	kstat
UNIXSYSCPU	API	5.2 perfstat_cpu 5.3 /usr/bin/mpstat	get_sysinfo	kstat
UNIXUSRCPU	API	5.2 perfstat_cpu 5.3 /usr/bin/mpstat	get_sysinfo	kstat
UNIXWAITIO	API	5.2 perfstat_cpu 5.3 /usr/bin/mpstat	get_sysinfo	kstat
UPTIME	API	time()-boottime()	time()-boottime()	time()-boottime()
VMFREEPRC	API	perfstat_ memory_ total	pstat_getdynamic+ /usr/sbin/swapinfo	sysconf +swapctl
VMFREEMEM	API	perfstat_ memory_ total	pstat_ getdynamic	sysconf
VMFREESWAP	API	perfstat_ memory_ total	/usr/sbin/swapinfo	swapctl
VMINPGWAIT	API	perfstat_cpu_ total	pstat_ getdynamic	kstat
VMINRUNQ	API	perfstat_cpu_ total	pstat_ getdynamic	kstat
VMPGFAULTS	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGIN	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGOUT	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGRCLM	API	perfstat_ memory_ total	pstat_ getvminfo	kstat

Table 86. Mechanisms used to gather UNIXOS / SP2OS attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
VMPGSIN	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMPGSOUT	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMSCAN	API	perfstat_ memory_ total	pstat_ getvminfo	kstat
VMUSEDPRC	API	perfstat_ memory_ total	pstat_ getdynamic+ /usr/sbin/swapinfo	sysconf +swapctl
XMTINT	API	perfstat_cpu_ total	get_sysinfo	kstat
ZID	API	N/A	N/A	zone_list
ZONE	API	N/A	N/A	zone_getattr

Table 87. Mechanisms used to gather UNIXPING attributes

Attribute	Collection methods	AIX	HPUX	Solaris
HOSTRESP	API	/usr/sbin/ping	/usr/sbin/ping	/usr/sbin/ping
PINGHOST	API	/usr/sbin/ping	/usr/sbin/ping	/usr/sbin/ping
PINGRESULT	API	/usr/sbin/ping	/usr/sbin/ping	/usr/sbin/ping

Table 88. Mechanisms used to gather UNXPRINTQ attributes

Attribute	Collection methods	AIX	HPUX	Solaris
DEVICENM	API	/bin/enq	N/A	N/A
PRINTQSIZE	API	/bin/enq	N/A	N/A
PRNTQDEPTH	API	/bin/enq	N/A	N/A
PRNTQSTATE	API	/bin/enq	N/A	N/A
PRT	API	/bin/enq	N/A	N/A

Table 89. Mechanisms used to gather UNIXPS attributes

Attribute	Collection methods	AIX	HPUX	Solaris
ADDR	API	getprocs64	pstat_getproc	/proc/%s/psinfo
BCMD	API	getprocs64	pstat_getproc	/proc/%s/psinfo
CHILDSTIME	API	getprocs64	pstat_getproc	/proc/%s/status
CHILDTIME	API	getprocs64	pstat_getproc	/proc/%s/status
CHILDUTIME	API	getprocs64	pstat_getproc	/proc/%s/status
CMD	API	getprocs64	pstat_getproc	/proc/%s/psinfo
COMMAND	API	getprocs64	pstat_getproc	/proc/%s/psinfo

Table 89. Mechanisms used to gather UNIXPS attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
CONTSWITCH	API	getprocs64	pstat_getproc	/proc/%s/usage
CPU	API	getprocs64	pstat_getproc	/proc/%s/psinfo
CPUID	API	getprocs64	pstat_getproc	/proc/%s/psinfo
CPUPERCENT	API	getprocs64	pstat_getproc	/proc/%s/psinfo
CPUTIME	API	getprocs64	pstat_getproc	/proc/%s/psinfo
EGID	API	getprocs64	pstat_getproc	/proc/%s/psinfo
EGRPN	API	getprocs64	pstat_getproc	/proc/%s/psinfo
ELAPTIME	API	getprocs64	pstat_getproc	/proc/%s/psinfo
EUID	API	getprocs64	pstat_getproc	/proc/%s/psinfo
EUSERN	API	getprocs64	pstat_getproc	/proc/%s/psinfo
EVENT	API	getprocs64	pstat_getproc	/proc/%s/psinfo
EXECSTATE	API	getprocs64	pstat_getproc	/proc/%s/psinfo
FLAG	API	getprocs64	pstat_getproc	/proc/%s/psinfo
GID	API	getprocs64	pstat_getproc	/proc/%s/psinfo
GRPN	API	getprocs64	pstat_getproc	/proc/%s/psinfo
HEAP	API	getprocs64	pstat_getproc	/proc/%s/psinfo
INVCONTSWT	API	getprocs64	pstat_getproc	/proc/%s/usage
MAJORFAULT	API	getprocs64	pstat_getproc	/proc/%s/usage
MEMPERCENT	API	getprocs64	pstat_getproc	/proc/%s/psinfo
MINORFAULT	API	getprocs64	pstat_getproc	/proc/%s/usage
NICE	API	getprocs64	pstat_getproc	/proc/%s/psinfo
PGID	API	getprocs64	pstat_getproc	/proc/%s/psinfo
PID	API	getprocs64	pstat_getproc	/proc/%s/psinfo
PPID	API	getprocs64	pstat_getproc	/proc/%s/psinfo
PRIORITY	API	getprocs64	pstat_getproc	/proc/%s/psinfo
PSU	API	getprocs64 (struct procentry64. pi_dvm)	N/A	N/A
RDS	API	getprocs64 (struct procentry64. pi_drss)	N/A	N/A
READWRITE	API	getprocs64	pstat_getproc	/proc/%s/psinfo
REPTYPE	API	getprocs64	pstat_getproc	/proc/%s/psinfo
RTS	API	getprocs64 (struct procentry64. pi_trss)	N/A	N/A
SAMPCPUPCT	API	getprocs64	pstat_getproc	/proc/%s/psinfo
SCHEDCLASS	API	getprocs64	pstat_getproc	/proc/%s/psinfo
SESSIONID	API	getprocs64	pstat_getproc	/proc/%s/psinfo
SIZE	API	getprocs64	pstat_getproc	/proc/%s/psinfo
STACK	API	getprocs64	pstat_getproc	/proc/%s/psinfo

Table 89. Mechanisms used to gather UNIXPS attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
STARTTIME	API	getprocs64	pstat_getproc	/proc/%s/psinfo
SYSTEMTIM	API	getprocs64	pstat_getproc	/proc/%s/status
SYSTEMTYPE	API	uname	uname	uname
THREADCNT	API	getprocs64	pstat_getproc	/proc/%s/status
TIME	API	getprocs64	pstat_getproc	/proc/%s/psinfo
TOTALTIME	API	getprocs64	pstat_getproc	/proc/%s/status
TOTCPUPERC	API	getprocs64	pstat_getproc	/proc/%s/psinfo
TTY	API	getprocs64	pstat_getproc	/proc/%s/psinfo
UID	API	getprocs64	pstat_getproc	/proc/%s/psinfo
UCMD	API	getprocs64	pstat_getproc	/proc/%s/psinfo
UCOMMAND	API	getprocs64	pstat_getproc	/proc/%s/psinfo
USERNAME	API	getprocs64	pstat_getproc	/proc/%s/psinfo
USERTIME	API	getprocs64	pstat_getproc	/proc/%s/status
UUSERNAME	API	getprocs64	pstat_getproc	/proc/%s/psinfo
VSIZE	API	getprocs64	pstat_getproc	/proc/%s/psinfo
WAITCPUTIM	API	getprocs64	pstat_getproc	/proc/%s/usage
WAITLKTIME	API	getprocs64	pstat_getproc	/proc/%s/usage
WLM_NAME	API	getprocs64 (struct procentry64. pi_classname)	N/A	N/A
WPAR_NAME	API	getprocs64 (struct procentry64. getcorralname (pi_cid))	N/A	N/A
ZONEID	API	N/A	N/A	/proc/%s/psinfo
ZONENAME	API	N/A	N/A	getzonename byid

Table 90. Mechanisms used to gather UNIXSOLZON attributes

Attribute	Collection methods	AIX	HPUX	Solaris
CAPCPU	API	prctl	prctl	prctl
CAPMEM	API	zonecfg	zonecfg	zonecfg
CPUSHARES	API	N/A	N/A	prctl
DEDCPU	API	zonecfg	zonecfg	zonecfg
IPID	API	N/A	N/A	zone_getattr
POOLID	API	N/A	N/A	zone_getattr
SCHED	API	N/A	N/A	pconf_info
SHAREPCT	API	N/A	N/A	prctl
ZCPU	API	N/A	N/A	prstat
ZCPUS	API	N/A	N/A	pconf_info

Table 90. Mechanisms used to gather UNIXSOLZON attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
ZID	API	N/A	N/A	zone_list
ZONENAME	API	N/A	N/A	zone_getattr
ZPATH	API	N/A	N/A	zone_getattr
ZRSS	API	N/A	N/A	prstat
ZSTATUS	API	N/A	N/A	zone_getattr
ZVMS	API	N/A	N/A	prstat

Table 91. Mechanisms used to gather UNIXTCP attributes

Attribute	Collection methods	AIX	HPUX	Solaris
PKTRETRPS	API	netstat -s -p tcp	get_mib_info	kstat

Table 92. Mechanisms used to gather UNIXUSER attributes

Attribute	Collection methods	AIX	HPUX	Solaris
UID	API	getpwuid_r	getpwuid_r	getpwuid_r
USERIDLE	API	stat(tty)	stat(tty)	stat(tty)
USERLOGIN	API	getutxent	getutxent	getutxent
USERNAME	API	getpwnam_r	getpwnam_r	getpwnam_r
USERSITE	API	getutxent	getutxent	getutxent
USERTTY	API	getutxent	getutxent	getutxent
USERWHEN	API	gmtime_r	gmtime_r	gmtime_r
UUSERLOGIN	API	getutxent	getutxent	getutxent
UUSERNAME	API	getpwnam_r	getpwnam_r	getpwnam_r

Table 93. Mechanisms used to gather UNIXWPARCP attributes

Attribute	Collection methods	AIX	HPUX	Solaris
CCL	AIX Script Data Provider	/usr/bin/iostat (RC_CPU_Limits _Hard_Max / 100)	N/A	N/A
LCCP	AIX Script Data Provider	/usr/bin/iostat (Num_CPUs _Consumed / LPAR_Entitlement)	N/A	N/A
LE	AIX Script Data Provider	/usr/bin/iostat -@ [*wpar_name*]	N/A	N/A
NCC	AIX Script Data Provider	/usr/bin/iostat -@ [*wpar_name*]	N/A	N/A
RCLHM	AIX Script Data Provider	/usr/sbin/lswpar -s A -qca	N/A	N/A

Table 93. Mechanisms used to gather UNIXWPARCP attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
SCP	AIX Script Data Provider	/usr/bin/iostat -@ [*wpar_name*]	N/A	N/A
UCP	AIX Script Data Provider	/usr/bin/iostat -@ [*wpar_name*]	N/A	N/A
WCCP	AIX Script Data Provider	/usr/bin/iostat -@ [*wpar_name*]	N/A	N/A
WPAR_NAME	AIX Script Data Provider	/usr/sbin/lswpar -s A -qca	N/A	N/A

Table 94. Mechanisms used to gather UNIXWPARFS attributes

Attribute	Collection methods	AIX	HPUX	Solaris
DN	AIX Script Data Provider	/usr/sbin/lswpar -cM -d,	N/A	N/A
MO	AIX Script Data Provider	/usr/sbin/lswpar -cM -d,	N/A	N/A
MP	AIX Script Data Provider	/usr/sbin/lswpar -cM -d,	N/A	N/A
NODE_NAME	AIX Script Data Provider	/usr/sbin/lswpar -cM -d,	N/A	N/A
VSF_TYPE	AIX Script Data Provider	/usr/sbin/lswpar -cM -d,	N/A	N/A
WPAR_NAME	AIX Script Data Provider	/usr/sbin/lswpar -cM -d,	N/A	N/A

Table 95. Mechanisms used to gather UNIXWPARIN attributes

Attribute	Collection methods	AIX	HPUX	Solaris
AO	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
API	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
AST	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A

Table 95. Mechanisms used to gather UNIXWPARIN attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
AUTOSTART	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
C	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
HOME	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
HOSTNAME	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
IP_ADDRESS	AIX Script Data Provider	Note: The perl code does not provide this attribute.	N/A	N/A
OWNER	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RCLHM	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RCLM	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RCLSM	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RCS	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RIA	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RMP	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RMT	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RMLHM	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RMLM	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RMLSM	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A

Table 95. Mechanisms used to gather UNIXWPARIN attributes (continued)

Attribute	Collection methods	AIX	HPUX	Solaris
RMS	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RPPVL	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
RC_RSET	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
SUD	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
STATE	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
TYPE	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
WAP	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A
WPAR_NAME	AIX Script Data Provider	/usr/sbin/lswpar -d, -qca [*wpar_name*]	N/A	N/A

Table 96. Mechanisms used to gather UNIXWPARNE attributes

Attribute	Collection methods	AIX	HPUX	Solaris
BI	AIX Script Data Provider	/usr/sbin/lswpar -d, -Nqa	N/A	N/A
IN	AIX Script Data Provider	/usr/sbin/lswpar -d, -Nqa	N/A	N/A
IP_ADDRESS	AIX Script Data Provider	/usr/sbin/lswpar -d, -Nqa	N/A	N/A
NM	AIX Script Data Provider	/usr/sbin/lswpar -d, -Nqa	N/A	N/A
WPAR_NAME	AIX Script Data Provider	/usr/sbin/lswpar -d, -Nqa	N/A	N/A

Table 97. Mechanisms used to gather UNIXWPARM attributes

Attribute	Collection methods	AIX	HPUX	Solaris
FMM	AIX Script Data Provider	/usr/bin/svmon -@ [wpar_name]	N/A	N/A
FMP	AIX Script Data Provider	FMM / (FMM + UMM) *100	N/A	N/A
LMSM	AIX Script Data Provider	/usr/bin/svmon	N/A	N/A
LMUP	AIX Script Data Provider	UMM / LMSM *100	N/A	N/A
MSM	AIX Script Data Provider	/usr/bin/svmon -@ [*wpar_name*]	N/A	N/A
RMLHM	AIX Script Data Provider	/usr/sbin/lswpar -qca	N/A	N/A
UMM	AIX Script Data Provider	/usr/bin/svmon -@ [*wpar_name*]	N/A	N/A
UMP	AIX Script Data Provider	UMM / (FMM + UMM) *100	N/A	N/A
WPAR_NAME	AIX Script Data Provider	/usr/sbin/lswpar -qca	N/A	N/A

Appendix F. Discovery Library Adapter for the monitoring agent

This chapter contains information about the Discovery Library Adapter (DLA) for the Monitoring Agent for UNIX.

About the DLA

The Tivoli Management Services DLA discovers resources and relationships and creates a Discovery Library Book file. The Book follows the Discovery Library IdML schema version 2.9.2 and is used to populate the Configuration Management Database (CMDB) and Tivoli Business System Management products. The Tivoli Management Services DLA discovers UNIX resources. For all UNIX systems that are active and online at the Tivoli Enterprise Portal Server, information is included in the discovery book for those resources. The Tivoli Management Services DLA discovers active resources. It is run on demand and can be run periodically to discover resources that were not active during previous discoveries.

The DLA discovers UNIX components.

More information about DLAs

The following sources contain additional information about using the DLA program with all monitoring agents:

- The *IBM Tivoli Monitoring Administrator's Guide* contains information about using the Tivoli Management Services Discovery Library Adapter.
- For information about using a DLA with Tivoli Application Dependency Discovery Manager (TADDM), see the information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/topic/com.ibm.taddm.doc_7.1/cmdb_welcome.html

UNIX data model class types represented in CDM

This section contains information about how the various source application data objects map to classes in the Common Data Model (CDM) for the Monitoring Agent for UNIX.

The following information is provided for each class where appropriate:

Relationships
CDM relationships (hierarchical) between currently identified model objects

CDM attributes, agent attributes, descriptions, and examples
CDM and agent attributes that are required to create an instance of a resource, descriptions of the attributes, and examples of the attributes

UNIX class

The following information describes the UNIX class.

CDM class name
sys.aix.Aix or sys.sun.Solaris or sys.hpux.HpUx

Relationships

- installedOn
- runsOn

CDM attributes, agent attributes, descriptions, and examples

- CDM attribute: ManagedSystemName
  Agent attribute: none
  Description: Managed System Name
- CDM attribute: OSVersion
  Agent attribute: SYSTEMVERS/UNIXOS
  Description: OS Version
- CDM attribute: Name
  Agent attribute: SYSTEMTYPE/UNIXOS
  Description: OS Type
- CDM attribute: FQDN
  Agent attribute: DNSNAME/UNIXIPADDR
  Description: Fully Qualified Domain Name

ComputerSystem class

The following information describes the ComputerSystem class.

CDM class name

sys.ComputerSystem

CDM attributes, agent attributes, descriptions, and examples

- CDM attribute: ManagedSystemName
  Agent attribute: none
  Description: Managed System Name
- CDM attribute: Name
  Agent attribute: none
  Description: Fully Qualified Host Name
- CDM attribute: Signature
  Agent attribute: IPADDRESS/UNIXIPADDR and MACADDRESS/UNIXNET
  Description: Lowest IP Address (MAC Address)
- CDM attribute: PrimaryMACAddress
  Agent attribute: MACADDRESS/UNIXNET
  Description: MAC Address of the network interface with the lowest IP Address (alpha order)
- CDM attribute: Type
  Agent attribute: none
  Description: "ComputerSystem"
- CDM attribute: Fqdn
  Agent attribute: DNSNAME/UNIXIPADDR
  Description: Fully Qualified Domain Name
- CDM attribute: SystemBoardUUID
  Agent attribute: UUID/UNIXMACHIN
  Description: System Board UUID
- CDM attribute: SerialNumber
  Agent attribute: MACSERIAL/UNIXMACHIN
  Description: Serial Number
- CDM attribute: Model
  Agent attribute: MODEL/UNIXMACHIN

Description: Model
- CDM attribute: Manufacturer
  Agent attribute: VENDOR/UNIXMACHIN
  Description: Manufacturer
- CDM attribute: VMID
  Agent attribute: VMID/UNIXMACHIN
  Description: Partition ID
- CDM attribute: Label
  Agent attribute: none
  Description: Fully Qualified Host Name

IpInterface class

The following information describes the IpInterface class.

CDM class name
net.IpInterface

Relationships

- contains

CDM attributes, agent attributes, descriptions, and examples
none

IpV4Address class

The following information describes the IpV4Address class.

CDM class name
net.IpV4Address

Relationships

- bindsTo

CDM attributes, agent attributes, descriptions, and examples

- CDM attribute: DotNotation
  Agent attribute: IPADDRESS/UNIXIPADDR
  Description: IP Address of the network interface
- CDM attribute: Label
  Description: IP Address of the network interface

IpV6Address class

The following information describes the IpV6Address class.

CDM class name
net.IpV6Address

Relationships

- bindsTo

CDM attributes, agent attributes, descriptions, and examples

- CDM attribute: StringNotation
  Agent attribute: IPADDRESS/LNXIPADDR
  Description: IP Address of the network interface
- CDM attribute: Label
  Description: IP Address of the network interface

Fqdn class

The following information describes the Fqdn class.

CDM class name

net.Fqdn

CDM attributes, agent attributes, descriptions, and examples

- CDM attribute: Fqdn
  Agent attribute: IPADDRESS/LNXIPADDR
  Description: Fully Qualified Domain Name for the network interface

TMSAgent class

The following information describes the TMSAgent class.

CDM class name

app.TMSAgent

Relationships

- installedOn
- monitors

CDM attributes, agent attributes, descriptions, and examples

- CDM attribute: ManagedSystemName
  Agent attribute: none
  Description: Managed System Name
- CDM attribute: ManagedObjectName
  Description: "p@" Managed System Name
- CDM attribute: SoftwareVersion
  Description: OS Agent ITM version
- CDM attribute: ProductCode
  Description: OS Agent Product Code (UX)
- CDM attribute: Affinity
  Description: OS Agent affinity
- CDM attribute: Label
  Description: Managed System Name "- UNIX OS"

Appendix G. Documentation library

This appendix contains information about the publications related to IBM Tivoli Monitoring and to the commonly shared components of Tivoli Management Services. These publications are listed in the following categories:

- IBM Tivoli Monitoring library
- Related publications

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* in the IBM Tivoli Monitoring and OMEGAMON XE Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/. To open the *Documentation Guide* in the information center, select **Using the publications** in the **Contents** pane.

To find a list of new and changed publications, click **What's new** on the Welcome page of the IBM Tivoli Monitoring and OMEGAMON XE Information Center. To find publications from the previous version of a product, click **Previous versions** under the name of the product in the **Contents** pane.

IBM Tivoli Monitoring library

The following publications provide information about IBM Tivoli Monitoring and about the commonly shared components of Tivoli Management Services:

- *Quick Start Guide*

  Introduces the components of IBM Tivoli Monitoring.
- *Installation and Setup Guide*, GC32-9407

  Provides instructions for installing and configuring IBM Tivoli Monitoring components on Windows, Linux, and UNIX systems.
- *Program Directory for IBM Tivoli Management Services on z/OS*, GI11-4105

  Gives instructions for the SMP/E installation of the Tivoli Management Services components on z/OS.
- *Configuring the Tivoli Enterprise Monitoring Server on z/OS*, SC27-2313

  Provides instructions for preparing, configuring, and customizing your monitoring servers on z/OS. This guide complements the *IBM Tivoli OMEGAMON XE and IBM Tivoli Management Services on z/OS Common Planning and Configuration Guide* and the *IBM Tivoli Monitoring Installation and Setup Guide*.
- *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.

- *High-Availability Guide for Distributed Systems*, SC23-9768

  Gives instructions for several methods of ensuring the availability of the IBM Tivoli Monitoring components.
- Tivoli Enterprise Portal online help

  Provides context-sensitive reference information about all features and customization options of the Tivoli Enterprise Portal. Also gives instructions for using and administering the Tivoli Enterprise Portal.
- *Tivoli Enterprise Portal User's Guide*, SC32-9409

  Complements the Tivoli Enterprise Portal online help. The guide provides hands-on lessons and detailed instructions for all Tivoli Enterprise Portal features.
- *Command Reference*, SC32-6045

  Provides detailed syntax and parameter information, as well as examples, for the commands you can use in IBM Tivoli Monitoring.
- *Troubleshooting Guide*, GC32-9458

  Provides information to help you troubleshoot problems with the software.
- *Messages*, SC23-7969

  Lists and explains messages generated by all IBM Tivoli Monitoring components and by z/OS-based Tivoli Management Services components (such as Tivoli Enterprise Monitoring Server on z/OS and TMS:Engine).
- *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.
- *IBM Tivoli Universal Agent API and Command Programming Reference Guide*, SC32-9461

  Explains the procedures for implementing the IBM Tivoli Universal Agent APIs and provides descriptions, syntax, and return status codes for the API calls and command-line interface commands.
- *Agent Builder User's Guide*, SC32-1921

  Explains how to use the Agent Builder for creating monitoring agents and their installation packages, and for adding functions to existing agents.
- *Performance Analyzer User's Guide*, SC27-4004

  Explains how to use the Performance Analyzer to understand resource consumption trends, identify problems, resolve problems more quickly, and predict and avoid future problems.

Documentation for the base agents

If you purchased IBM Tivoli Monitoring as a product, you received a set of *base* monitoring agents as part of the product. If you purchased a monitoring agent product (for example, an OMEGAMON XE product) that includes the commonly shared components of Tivoli Management Services, you did not receive the base agents.

The following publications provide information about using the base agents.

- Operating system agents:
  - *Windows OS Agent User's Guide*, SC32-9445
  - *UNIX OS Agent User's Guide*, SC32-9446
  - *Linux OS Agent User's Guide*, SC32-9447
  - *i5/OS Agent User's Guide*, SC32-9448
  - *UNIX Log Agent User's Guide*, SC32-9471
- Agentless operating system monitors:
  - *Agentless Monitoring for Windows Operating Systems User's Guide*, SC23-9765
  - *Agentless Monitoring for AIX Operating Systems User's Guide*, SC23-9761
  - *Agentless Monitoring for HP-UX Operating Systems User's Guide*, SC23-9763
  - *Agentless Monitoring for Solaris Operating Systems User's Guide*, SC23-9764
  - *Agentless Monitoring for Linux Operating Systems User's Guide*, SC23-9762
- Warehouse agents:
  - *Warehouse Summarization and Pruning Agent User's Guide*, SC23-9767
  - *Warehouse Proxy Agent User's Guide*, SC23-9766
- System P agents:
  - *AIX Premium Agent User's Guide*, SA23-2237
  - *CEC Base Agent User's Guide*, SC23-5239
  - *HMC Base Agent User's Guide*, SA23-2239
  - *VIOS Premium Agent User's Guide*, SA23-2238
- Other base agents:
  - *Systems Director base Agent User's Guide*, SC27-2872
  - *Tivoli Log File Agent User's Guide*, SC14-7484
  - *Monitoring Agent for IBM Tivoli Monitoring 5.x Endpoint User's Guide*, SC32-9490

Related publications

You can find useful information about related products in the IBM Tivoli Monitoring and OMEGAMON XE Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/.

Other sources of documentation

You can also obtain technical documentation about IBM Tivoli Monitoring and related products from the following sources:

- IBM Integrated Service Management Library

  http://www-01.ibm.com/software/brandcatalog/ismlibrary/

  IBM Integrated Service Management Library is an online catalog that contains integration documentation and other downloadable product extensions.
- Redbooks

  http://www.redbooks.ibm.com/

  IBM Redbooks® and Redpapers include information about products from platform and solution perspectives.

- Technotes

  Technotes provide the latest information about known product limitations and workarounds. You can find Technotes through the IBM Software Support Web site at http://www.ibm.com/software/support.
- Tivoli wikis on the IBM developerWorks Web site

  Tivoli Wiki Central at http://www.ibm.com/developerworks/wikis/display/tivoli/Home is the home for interactive wikis that offer best practices and scenarios for using Tivoli products. The wikis contain white papers contributed by IBM employees, and content created by customers and business partners.

  Two of these wikis are of particular relevance to IBM Tivoli Monitoring:

  – Tivoli Distributed Monitoring and Application Management Wiki at http://www.ibm.com/developerworks/wikis/display/tivolimonitoring/Home provides information about IBM Tivoli Monitoring and related distributed products, including IBM Tivoli Composite Application Management products.
  – Tivoli System z Monitoring and Application Management Wiki at http://www.ibm.com/developerworks/wikis/display/tivoliomegamon/Home provides information about the OMEGAMON XE products, NetView for z/OS, Tivoli Monitoring Agent for z/TPF, and other System z monitoring and application management products.

Appendix H. Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in this product enable users to do the following:

- Use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. Consult the product documentation of the assistive technology for details on using those technologies with this product.
- Operate specific or equivalent features using only the keyboard.
- Magnify what is displayed on the screen.

In addition, the product documentation was modified to include the following features to aid accessibility:

- All documentation is available in both HTML and convertible PDF formats to give the maximum opportunity for users to apply screen-reader software.
- All images in the documentation are provided with alternative text so that users with vision impairments can understand the contents of the images.

Navigating the interface using the keyboard

Standard shortcut and accelerator keys are used by the product and are documented by the operating system. Refer to the documentation provided by your operating system for more information.

Magnifying what is displayed on the screen

You can enlarge information on the product windows using facilities provided by the operating systems on which the product is run. For example, in a Microsoft Windows environment, you can lower the resolution of the screen to enlarge the font sizes of the text on the screen. Refer to the documentation provided by your operating system for more information.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 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 be displayed.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

Index

A

accessibility 331
actions
  See Take Action commands
agent
  trace logs 201
  troubleshooting 210
agent installation problems 206
AIX AMS attribute group
  UNIXAMS historical table 275
AIX Defined Users attribute group
  UNIXDUSERS historical table 278
AIX Devices attribute group
  UNIXDEVIC historical table 277
AIX LPAR attribute group
  UNIXLPAR historical table 280
AIX WPAR attribute group
  UNIXWPARIN historical table 293
AIX WPAR CPU attribute group
  UNIXWPARCP historical table 292
AIX WPAR File System attribute group
  UNIXWPARFS historical table 292
AIX WPAR Network attribute group
  UNIXWPARNE historical table 294
AIX WPAR Physical Memory attribute group
  UNIXWPARPM historical table 294
All Users attribute group
  UNIXALLUSR historical table 275
AMS Recycle Agent Instance action 132
AMS Reset Agent Daily Restart Count action 132
AMS Start Agent action 133
AMS Start Agent Instance action 133
AMS Start Management action 134
AMS Stop Agent action 134
AMS Stop Management action 135
attribute
  cross-reference 36
  multiple-instance definition 33
  single-instance definition 33
attribute groups
  AIX AMS
    UNIXAMS historical table 275
  AIX Defined Users
    UNIXDUSERS historical table 278
  AIX Devices
    UNIXDEVIC historical table 277
  AIX LPAR
    UNIXLPAR historical table 280
  AIX WPAR
    UNIXWPARIN historical table 293
  AIX WPAR CPU
    UNIXWPARCP historical table 292
  AIX WPAR File System
    UNIXWPARFS historical table 292
  AIX WPAR Network
    UNIXWPARNE historical table 294
  AIX WPAR Physical Memory
    UNIXWPARPM historical table 294
  All Users
    UNIXALLUSR historical table 275
attribute groups *(continued)*
  Disk Information
    UNIXDISK historical table 277
  Disk Performance
    UNIXDPERF historical table 278
  File Information
    UNIXFILE historical table 278
  Group
    UNIXGROUP historical table 279
  IP Address
    UNIXIPADDR historical table 279
  Machine Information
    UNIXMACHIN historical table 281
  more information 34
  Network
    UNIXNET historical table 283
  NFS and RPC Statistics
    UNIXNFS historical table 283
  overview 33
  Ping
    UNIXPING historical table 288
  Print Queue
    UNIXPRINTQ historical table 288
  Process
    UNIXPS historical table 289
  SMP CPU
    UNIXCPU historical table 276
  Solaris Zones
    UNIXSOLZON historical table 291
  System
    UNIXOS historical table 285
  TCP Statistics
    UNIXTCP historical table 291
  UNIX Memory
    UNIXMEM historical table 281
  User
    UNIXUSER historical table 291
attributes
  more information 34
  overview 33

B

books
  see publications 223
built-in troubleshooting features 199

C

calculate historical data disk space 115
capacity planning for historical data 115
cluster environment 11
code, product 5
commands, Take Action 131
components 5

D

data
  trace logs 200

data collection 297
data provider
See agent
database agent installation problems 206
developerWorks Web site 330
disk capacity planning for historical data 115
Disk Information attribute group
UNIXDISK historical table 277
Disk Performance attribute group
UNIXDPERF historical table 278
disk space requirements 7
documentation
See publications

E

education
see Tivoli technical training 223
environment
features 1
event
mapping 245

F

features, Monitoring Agent for UNIX OS 1
File Information attribute group
UNIXFILE historical table 278
files
agent trace 201
installation trace 201
other trace log 202
trace logs 200

G

gathering support information 199
Group attribute group
UNIXGROUP historical table 279

H

historical data
calculate disk space 115
disk capacity planning 115
historical reports
attributes 36
column header 36
cross reference 36
historical tables
UNIXALLUSR 275
UNIXAMS 275
UNIXCPU 276
UNIXDEVIC 277
UNIXDISK 277
UNIXDPERF 278
UNIXDUSERS 278
UNIXFILE 278
UNIXGROUP 279
UNIXIPADDR 279
UNIXLPAR 280
UNIXMACHIN 281
UNIXMEM 281
UNIXNET 283
UNIXNFS 283
historical tables *(continued)*
UNIXOS 285
UNIXPING 288
UNIXPRINTQ 288
UNIXPS 289
UNIXSOLZON 291
UNIXTCP 291
UNIXUSER 291
UNIXWPARCP 292
UNIXWPARFS 292
UNIXWPARIN 293
UNIXWPARNE 294
UNIXWPARPM 294

I

IBM Software Support
See support
IBM Support Assistant 222
IBM Tivoli Enterprise Console
event mapping 245
optional product 6
IBM Tivoli Monitoring for UNIX
benefits of 1
definition of 1
types of information collected by 1
information, additional
attributes 34
policies 137
situations 118
Take Action commands 131
workspaces 13
installation
log file 201
problems 206
installation requirements 7
Integrated Service Management Library documentation 329
interface, user 6
troubleshooting for Tivoli Enterprise Portal 215
IP Address attribute group
UNIXIPADDR historical table 279
ISA 222

L

libraries
IBM Tivoli Monitoring 327
limited user permissions, upgrading your warehouse with 226
logging
agent trace logs 201, 202
built-in features 199
installation log files 201
location and configuration of logs 200
trace log files 200

M

Machine Information attribute group
UNIXMACHIN historical table 281
manuals
see publications 223
memory requirements 7
messages
built-in features 199

Monitoring Agent for UNIX OS
components 5
features 1
Monitoring Agent for UNIX OS installation problems 206
multiple-instance attributes
definition 33

N

Network attribute group
UNIXNET historical table 283
NFS and RPC Statistics attribute group
UNIXNFS historical table 283
nonroot users 9

O

online publications
accessing 223
operating systems 7
ordering publications 223
other requirements 8

P

path names, for trace logs 200
permissions, upgrading your warehouse with limited user 226
Ping attribute group
UNIXPING historical table 288
policies
list of all 137
more information 137
overview 137
predefined 137
UNIX_CPU_Busy 137
UNIX_Disk_Space_Full 138
UNIX_Virtual_Memory_High 138
Print Queue attribute group
UNIXPRINTQ historical table 288
problems and workarounds 205
problems with monitoring UNIX 221
Process attribute group
UNIXPS historical table 289
product code 5
publications
accessing online 223
developerWorks Web site 330
OPAL
ISM 329
ordering 223
Redbooks 329
related 329
Technotes 330
types 327
wikis 330
purposes
troubleshooting 199

Q

queries, using attributes 33

R

Redbooks 329
remote deployment
troubleshooting 216
requirements
disk space 7
memory 7
operating system 7
other 8
requirements, installation 7

S

Sample_kill_Process Take Action command 135
security 9
single-instance attributes
definition 33
situations
general troubleshooting 219
list of all 118
more information 118
overview 117
predefined 118
specific troubleshooting 218
UMX_AMS_Alert_Critical 120
UNIX_Active_Virtual_Memory 125
UNIX_AIX_Avg_ReqInWaitQ_MS_Info 120
UNIX_AIX_Avg_Transfer_MS_Info 120
UNIX_AIX_CPU_CtxSwitch_Hi_Info 125
UNIX_AIX_Device_Stopped_Warning 125
UNIX_AIX_Memory_RePg_Hi_Info 125
UNIX_AIX_NetBandwidth_High_Info 122
UNIX_AIX_Process_ResDat_Hi_Info 123
UNIX_AIX_Process_ResTxt_Hi_Info 123
UNIX_AIX_ServQ_Full_PerSec_Info 120
UNIX_AIX_System_HypPct_Hi_Info 125
UNIX_AIX_System_NProcs_Hi_Info 126
UNIX_AIX_User_Acct_Locked_Info 126
UNIX_AIX_User_Login_Retry_Info 126
UNIX_BP_AvgCpuBusyPct5min_Criti 126
UNIX_BP_CpuBusyPct_Critical 126
UNIX_BP_LoadAvg5min_Critical 126
UNIX_BP_NetInOutErrPct_Critical 122
UNIX_BP_NumberZombies_Warning 127
UNIX_BP_PagingRate_Critical 127
UNIX_BP_ProcHighCpu_Critical 124
UNIX_BP_ProcMissing_Critical 124
UNIX_BP_SpaceUsedPct_Critical 121
UNIX_BP_SpaceUsedPctCustom_Crit 121
UNIX_BP_SwapSpaceUsedPct_Critic 127
UNIX_BP_SysWaitIOPct_Warning 127
UNIX_CMD_Disk_Inodes_Critical 121
UNIX_CMD_Disk_Inodes_Critical_2 121
UNIX_CMD_Disk_Space_Warning 121
UNIX_CMD_Process_Critical 124
UNIX_CMD_Runaway_Process 124
UNIX_CPU_Busy_Critical 127
UNIX_CPU_Busy_Warning 127
UNIX_CPU_Critical 124
UNIX_CPU_Warning 124
UNIX_Disk_Availability 121
UNIX_Filemount_Critical 121
UNIX_HD_Config_Critical 122
UNIX_HD_Config_Critical_2 122
UNIX_HD_Excessive_IO_Wait 127
UNIX_LPAR_MaxCPUCapUsed_Info 128

situations *(continued)*
UNIX_LPAR_Moved_Info 128
UNIX_LPARBusy_pct_Warning 128
UNIX_LPARentused_Info 128
UNIX_LPARfreepool_Warning 128
UNIX_LPARPhanIntrs_Info 128
UNIX_LPARPhyBusy_pct_Warning 128
UNIX_LPARvcs_Info 128
UNIX_Network_Collsns_Critical 122
UNIX_Network_Collsns_Warning 123
UNIX_Network_Errors 123
UNIX_Network_Interface_Busy 123
UNIX_Network_Interface_Idle 123
UNIX_NFS_RPC_Rejects 123
UNIX_Process_Memory_Critical 124
UNIX_Process_Memory_Leak 125
UNIX_Process_Memory_Warning 125
UNIX_Process_MISSING_inetd 125
UNIX_scratch_tmp_Disk_Full 122
UNIX_System_Busy_Critical 129
UNIX_System_Busy_Warning 129
UNIX_System_Capacity_Critical 129
UNIX_System_Paging_Critical 129
UNIX_System_Virtual_Memory_Warning 129
UNIX_User_CPU_Critical 129
UNIX_User_File_Exists 122
UNIX_WPAR_Admin_Op_Info 129
UNIX_WPAR_Broken_Warning 129
UNIX_WPAR_CPU_Usage_Warning 129
UNIX_WPAR_Mem_Usage_Warning 130
UNIX_WPAR_Min_CPU_Limit_Info 130
UNIX_WPAR_Min_Mem_Limit_Info 130
UNIX_WPAR_RC_Inactive_Info 130
UNIX_WPAR_Unlim_CPU_Shares_Info 130
UNIX_WPAR_Unlim_Mem_Shares_Info 130
situations, using attributes 33
SMP CPU attribute group
UNIXCPU historical table 276
Software Support 222
Solaris Zones attribute group
UNIXSOLZON historical table 291
support
gathering information for 199
support assistant 222
System attribute group
UNIXOS historical table 285

T

Take Action commands
AMS Recycle Agent Instance 132
AMS Reset Agent Daily Restart Count 132
AMS Start Agent 133
AMS Start Agent Instance 133
AMS Start Management 134
AMS Stop Agent 134
AMS Stop Management 135
more information 131
overview 131
Sample_kill_Process 135
troubleshooting 221
target application
problems 221
TCP Statistics attribute group
UNIXTCP historical table 291
Technotes 330
Terminal 292
Tivoli Data Warehouse 6
Tivoli Enterprise Console
See IBM Tivoli Enterprise Console
Tivoli Enterprise Monitoring Server 6
Tivoli Enterprise Portal
component 5
troubleshooting 215
Tivoli Information Center 223
Tivoli technical training 223
Tivoli user groups 223
trace logs 200
directories 200
training, Tivoli technical 223
transport command 221
troubleshooting 199, 205
agents 210
built-in features 199
installation 206
installation logs 201
remote deployment 216
situations 217, 219
Take Action commands 221
Tivoli Enterprise Portal 215
uninstallation 206
uninstallation logs 201
workspaces 216

U

UMX_AMS_Alert_Critical situation 120
uninstallation
log file 201
problems 206
UNIX Memory attribute group
UNIXMEM historical table 281
UNIX problems 221
UNIX_Active_Virtual_Memory situation 125
UNIX_AIX_Avg_ReqInWaitQ_MS_Info situation 120
UNIX_AIX_Avg_Transfer_MS_Info situation 120
UNIX_AIX_CPU_CtxSwitch_Hi_Info situation 125
UNIX_AIX_Device_Stopped_Warning situation 125
UNIX_AIX_Memory_RePg_Hi_Info situation 125
UNIX_AIX_NetBandwidth_High_Info situation 122
UNIX_AIX_Process_ResDat_Hi_Info situation 123
UNIX_AIX_Process_ResTxt_Hi_Info situation 123
UNIX_AIX_ServQ_Full_PerSec_Info situation 120
UNIX_AIX_System_HypPct_Hi_Info situation 125
UNIX_AIX_System_NProcs_Hi_Info situation 126
UNIX_AIX_User_Acct_Locked_Info situation 126
UNIX_AIX_User_Login_Retry_Info situation 126
UNIX_BP_AvgCpuBusyPct5min_Criti situation 126
UNIX_BP_CpuBusyPct_Critical situation 126
UNIX_BP_LoadAvg5min_Critical situation 126
UNIX_BP_NetInOutErrPct_Critical situation 122
UNIX_BP_NumberZombies_Warning situation 127
UNIX_BP_PagingRate_Critical situation 127
UNIX_BP_ProcHighCpu_Critical situation 124
UNIX_BP_ProcMissing_Critical situation 124
UNIX_BP_SpaceUsedPct_Critical situation 121
UNIX_BP_SpaceUsedPctCustom_Crit situation 121
UNIX_BP_SwapSpaceUsedPct_Critic situation 127
UNIX_BP_SysWaitIOPct_Warning situation 127
UNIX_CMD_Disk_Inodes_Critical situation 121
UNIX_CMD_Disk_Inodes_Critical_2 situation 121
UNIX_CMD_Disk_Space_Warning situation 121
UNIX_CMD_Process_Critical situation 124
UNIX_CMD_Runaway_Process situation 124

UNIX_CPU_Busy policy 137
UNIX_CPU_Busy_Critical situation 127
UNIX_CPU_Busy_Warning situation 127
UNIX_CPU_Critical situation 124
UNIX_CPU_Warning situation 124
UNIX_Disk_Availability situation 121
UNIX_Disk_Space_Full policy 138
UNIX_Filemount_Critical situation 121
UNIX_HD_Config_Critical situation 122
UNIX_HD_Config_Critical_2 situation 122
UNIX_HD_Excessive_IO_Wait situation 127
UNIX_LPAR_MaxCPUCapUsed_Info situation 128
UNIX_LPAR_Moved_Info situation 128
UNIX_LPARBusy_pct_Warning situation 128
UNIX_LPARentused_Info situation 128
UNIX_LPARfreepool_Warning situation 128
UNIX_LPARPhanIntrs_Info situation 128
UNIX_LPARPhyBusy_pct_Warning situation 128
UNIX_LPARvcs_Info situation 128
UNIX_Network_Collsns_Critical situation 122
UNIX_Network_Collsns_Warning situation 123
UNIX_Network_Errors situation 123
UNIX_Network_Interface_Busy situation 123
UNIX_Network_Interface_Idle situation 123
UNIX_NFS_RPC_Rejects situation 123
UNIX_Process_Memory_Critical situation 124
UNIX_Process_Memory_Leak situation 125
UNIX_Process_Memory_Warning situation 125
UNIX_Process_MISSING_inetd situation 125
UNIX_scratch_tmp_Disk_Full situation 122
UNIX_System_Busy_Critical situation 129
UNIX_System_Busy_Warning situation 129
UNIX_System_Capacity_Critical situation 129
UNIX_System_Paging_Critical situation 129
UNIX_System_Virtual_Memory_Warning situation 129
UNIX_User_CPU_Critical situation 129
UNIX_User_File_Exists situation 122
UNIX_Virtual_Memory_High policy 138
UNIX_WPAR_Admin_Op_Info situation 129
UNIX_WPAR_Broken_Warning situation 129
UNIX_WPAR_CPU_Usage_Warning situation 129
UNIX_WPAR_Mem_Usage_Warning situation 130
UNIX_WPAR_Min_CPU_Limit_Info situation 130
UNIX_WPAR_Min_Mem_Limit_Info situation 130
UNIX_WPAR_RC_Inactive_Info situation 130
UNIX_WPAR_Unlim_CPU_Shares_Info situation 130
UNIX_WPAR_Unlim_Mem_Shares_Info situation 130
UNIXALLUSR historical table 275
UNIXAMS historical table 275
UNIXCPU historical table 276
UNIXDEVIC historical table 277
UNIXDISK historical table 277
UNIXDPERF historical table 278
UNIXDUSERS historical table 278
UNIXFILE historical table 278
UNIXGROUP historical table 279
UNIXIPADDR historical table 279
UNIXLPAR historical table 280
UNIXMACHIN historical table 281
UNIXMEM historical table 281
UNIXNET historical table 283
UNIXNFS historical table 283
UNIXOS historical table 285
UNIXPING historical table 288
UNIXPRINTQ historical table 288
UNIXPS historical table 289
UNIXSOLZON historical table 291
UNIXTCP historical table 291
UNIXUSER historical table 291
UNIXWPARCP historical table 292
UNIXWPARFS historical table 292
UNIXWPARIN historical table 293
UNIXWPARNE historical table 294
UNIXWPARPM historical table 294
upgrading for warehouse summarization 225
upgrading your warehouse with limited user permissions 226
User attribute group
  UNIXUSER historical table 291
user groups, Tivoli 223
user interfaces options 6
user permissions, upgrading your warehouse with limited 226

W

Warehouse Proxy agent 6
warehouse summarization
  upgrading for
    overview 225
Warehouse Summarization and Pruning agent 6
warehouse summarization upgrading
  affected attribute groups and supporting scripts 229
  DB2 warehouse database procedure 230
  effects on summarized attributes 225
  MS SQL warehouse database procedure 232
  Oracle warehouse database procedure 231
  procedures for running scripts 230
  table summary 228
  tables in the warehouse 225
  types of table changes 227
  upgrading your warehouse 228
wikis 330
Windows agent installation problems 206
workarounds 205
  agents 210
  remote deployment 216
  situations 217
  Take Action commands 221
  Tivoli Enterprise Portal 215
  workspaces 216
workspaces
  list of all 13
  more information 13
  overview 13
  predefined 13
  troubleshooting 216

Printed in USA

SC32-9446-05