

IBM Tivoli Monitoring
Version 6.3

Linux OS Agent User's Guide



IBM Tivoli Monitoring
Version 6.3

Linux OS Agent User's Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 315.

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

© **Copyright IBM Corporation 1994, 2013.**

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

Contents

Tables	vii
-------------------------	------------

Chapter 1. Using the monitoring agent **1**

New in this release	1
Using IBM Tivoli Monitoring	3
Components of the monitoring agent	3
User interface options	4

Chapter 2. Requirements for the monitoring agent. **5**

Naming instances.	7
Running as a non-Administrator user	7
Using Agent Management Services	9
Filtering capabilities on the names of processes	10

Chapter 3. Workspaces. **11**

Organization of the predefined workspaces.	11
Agent Management Services workspace	12
Agents' Management Log workspace	12
All Files workspace.	13
Capacity Usage Information workspace	13
CPU Averages workspace.	14
Disk I/O Extended Rate workspace	14
Disk I/O Rate workspace.	14
Disk Usage workspace.	15
File Information workspace	15
Historical Summarized Availability workspace.	15
Historical Summarized Availability Daily workspace	15
Historical Summarized Availability Hourly workspace.	16
Historical Summarized Availability Weekly workspace.	16
Historical Summarized Capacity workspace	16
Historical Summarized Capacity Daily workspace	17
Historical Summarized Capacity Hourly workspace	17
Historical Summarized Capacity Weekly workspace	17
Historical Summarized Performance workspace	18
Historical Summarized Performance Daily workspace.	18
Historical Summarized Performance Hourly workspace.	18
Historical Summarized Performance Weekly workspace.	19
Linux workspace	19
Network workspace	19
NFS Statistics workspace	20
Process workspace	20
Process CPU Usage workspace	20
Process User Information workspace	21
RPC Statistics workspace	21
Sockets Information workspace.	21
Specific File Information workspace	22
System Configuration workspace	22
System Information workspace	22
Users Workspace	23

Virtual Memory Statistics workspace	23
Virtual Memory Usage Trends workspace	24

Chapter 4. Attributes. **25**

Agent Availability Management Status attributes	27
Agent Active Runtime Status attributes	28
Alerts Table attributes	29
All Users attributes.	30
Configuration Information attributes	31
CPU attributes	32
CPU attributes (superseded).	33
CPU Averages attributes	34
CPU Averages attributes (superseded)	36
CPU Configuration attributes	37
Disk attributes	38
Disk attributes (superseded).	39
Disk IO attributes	41
Disk IO attributes (superseded).	42
Disk Usage Trends attributes	43
Disk Usage Trends attributes (superseded)	44
File Comparison Group attributes	46
File Information attributes	46
File Pattern Group attributes	48
I/O Ext attributes	48
I/O Ext attributes (superseded).	50
IP Address attributes	51
Linux Group attributes	52
Linux Host Availability attributes	52
Linux TCP Statistics attributes	53
LPAR attributes	54
Machine Information attributes.	55
Network attributes	56
Network attributes (superseded)	60
NFS Statistics attributes	63
NFS Statistics attributes (superseded).	68
OS Configuration attributes	72
Process attributes	73
Process attributes (superseded).	78
Process User Info attributes	82
Process User Info attributes (superseded)	84
RPC Statistics attributes	86
RPC Statistics attributes (superseded).	87
Sockets Detail attributes	88
Sockets Detail attributes (superseded)	90
Sockets Status attributes	91
Sockets Status attributes (superseded)	91
Swap Rate attributes	92
Swap Rate attributes (superseded).	93
System Statistics attributes	94
System Statistics attributes (superseded).	96
User Login attributes	98
User Login attributes (superseded)	98
VM Stats attributes	99
VM Stats attributes (superseded)	102
Disk capacity planning for historical data	103

Chapter 5. Situations 107

Predefined situations	107
Linux_AMS_Alert_Critical situation	109
Linux_BP_AvgCpuBusyPct1h_Critic situation.	109
Linux_BP_CpuBusyPct_Critical situation	109
Linux_BP_CpuWaitIOPct_Warning situation	110
Linux_BP_LoadAvg5min_Critical situation.	110
Linux_BP_NetTotalErrPct_Warning situation	110
Linux_BP_NumberZombies_Warning situation	110
Linux_BP_ProcHighCpu_Critical situation	110
Linux_BP_ProcMissing_Critical situation	111
Linux_BP_SpaceUsedPct_Critical situation	111
Linux_BP_SpaceUsedPctCustom_Cri situation	111
Linux_BP_SwapSpaceUsedPct_Criti situation	111
Linux_Fragmented_File_System situation	111
Linux_Fragmented_File_System_2 situation	112
Linux_High_CPU_Overload situation	112
Linux_High_CPU_Overload_2 situation	112
Linux_High_CPU_System situation	112
Linux_High_CPU_System_2 situation	112
Linux_High_Packet_Collisions situation	112
Linux_High_Packet_Collisions_2 situation	113
Linux_High_RPC_Retransmit situation	113
Linux_High_RPC_Retransmit_2 situation	113
Linux_High_Zombies situation	113
Linux_High_Zombies_2 situation	113
Linux_Low_Pct_Inodes situation	113
Linux_Low_Pct_Inodes_2 situation	114
Linux_Low_percent_space situation	114
Linux_Low_percent_space_2 situation	114
Linux_Low_Space_Available situation	114
Linux_Low_Space_Available_2 situation	114
Linux_Network_Status situation	114
Linux_Network_Status_2 situation	115
Linux_NFS_Buffer_High situation	115
Linux_NFS_Buffer_High_2 situation	115
Linux_NFS_Getattr_High situation	115
Linux_NFS_Getattr_High_2 situation	115
Linux_NFS_rmlink_high situation	115
Linux_NFS_rmlink_high_2 situation	116
Linux_NFS_Read_High situation	116
Linux_NFS_Read_High_2 situation	116
Linux_NFS_Writes_High situation	116
Linux_NFS_Writes_High_2 situation.	116
Linux_Packets_Error situation	116
Linux_Packets_Error_2 situation	117
Linux_Process_High_Cpu situation	117
Linux_Process_High_Cpu_2 situation	117
Linux_Process_High_Instant_CPU situation	117
Linux_Process_stopped situation	117
Linux_Process_stopped_2 situation	117
Linux_RPC_Bad_Calls situation	118
Linux_RPC_Bad_Calls_2 situation	118
Linux_System_Thrashing situation	118
Linux_System_Thrashing_2 situation	118

Chapter 6. Take Action commands 119

Predefined Take Action commands	119
AMS Recycle Agent Instance	120
AMS Reset Agent Daily Restart Count	120
AMS Start Agent action	121

AMS Start Agent Instance action	121
AMS Stop Agent action	122
AMS Start Management action	122
AMS Stop Management action	123
Sample_kill_Process action	123

Chapter 7. Policies 125

Predefined policies	125
-------------------------------	-----

Chapter 8. Server dashboards 127

Server dashboards background information	127
Checking the health of your monitored environment.	130
Displaying situation event results	131
Managed System Groups Overview dashboard	132
Managed system group dashboard	133
Situation Events dashboard.	135
Situation event results dashboard.	136
Linux managed system dashboard	138
Page layout and controls	140
Table controls	141
Copying the URL	142
Launching to the Tivoli Enterprise Portal	142
Setting a trace	143

Chapter 9. Tivoli Common Reporting for the monitoring agent 145

Utilization Details for Single Resource report	150
Utilization Details for Multiple Resources report	153
Utilization Comparison for Single Resource report	156
Utilization Comparison for Multiple Resources report	158
Utilization Heat Chart for Single Resource report	162
Memory Utilization for Single Resource report	165
Memory Utilization for Multiple Resources Comparison report	167
Top Resources Utilization report	169
Top Situations by Status report	173
Enterprise Resources List report	174
Enterprise Daily Utilization Heat Chart report	174
Enterprise Summary report.	176
Top Resources by Availability	178
Top Resources Utilization Summary Heat Chart report	180
Top Resources by Availability (MTTR/MTBSI)	182
Resource Availability Comparison	184
Availability Heat Chart for Single Resource	186
CPU Utilization Comparison for Multiple Resources	188
CPU Utilization for Single Resource	191
Disk Utilization for Single Resource	192
Disk Utilization Comparison for Multiple Resources	195
Situations History report	198
Creating custom queries and reports	200

Chapter 10. Troubleshooting 207

Gathering product information for IBM Software Support	207
Built-in troubleshooting features	207

Problem classification	208
Trace logging	208
Principal trace log files	209
Setting RAS trace parameters	211
Problems and workarounds	212
Installation and configuration troubleshooting	212
Agent troubleshooting	217
Tivoli Enterprise Portal troubleshooting	219
Troubleshooting for remote deployment	220
Tivoli Common Reporting troubleshooting	220
Situation troubleshooting	221

Appendix A. Upgrading for warehouse summarization 227

Tables in the warehouse	227
Effects on summarized attributes	227
Upgrading your warehouse with limited user permissions	228
Types of table changes	229
Table summary	230
Upgrading your warehouse for primary key and tablespace changes	231
Affected attribute groups and supporting scripts	232
Procedures	232

Appendix B. IBM Tivoli Enterprise Console event mapping 237

Appendix C. Monitoring Agent for Linux OS data collection 263

KLZCPU attributes	263
KLZCPUAVG attributes	264
KLZDISK attributes	265
KLZDSKIO attributes	265
KLZDU attributes	266
KLZIOEXT attributes	266
KLZLOGIN attributes	267
KLZLPAR attributes	268
KLZNET attributes	268
KLZNFS attributes	269
KLZPROC attributes	271
KLZPUSR attributes	275
KLZRPC attributes	276
KLZSOCKD attributes	276
KLZSOCKS attributes	277
KLZSWPRT attributes	277
KLZSYS attributes	277
KLZVM attributes	278

LNXALLUSR attributes	280
LNXCPU attributes	280
LNXCPUAVG attributes	280
LNXCPUCON attributes	281
LNXDISK attributes	282
LNXDSKIO attributes	282
LNXDU attributes	283
LNXFILCMP attributes	283
LNXFILE attributes	284
LNXFILPAT attributes	284
LNXGROUP attributes	285
LNXIOEXT attributes	285
LNXIPADDR attributes	286
LNXLOGIN attributes	286
LNXMACHIN attributes	287
LNXNET attributes	287
LNXNFS attributes	288
LNXOSCON attributes	290
LNXPING attributes	290
LNXPROC attributes	291
LNXPUSR attributes	292
LNXRPC attributes	293
LNXSOCKD attributes	294
LNXSOCKS attributes	294
LNXSWPRT attributes	295
LNXSYS attributes	295
LNXVM attributes	296

Appendix D. Discovery Library Adapter for the monitoring agent . . . 297

Linux data model class types represented in CDM	297
-----------------------------------------------------------	-----

Appendix E. OSLC resources for the Linux OS agent 301

Documentation library 307	
IBM Tivoli Monitoring library	307
Documentation for the base agents	308
Related publications	309
Other sources of documentation	309

Support information 311

Notices 315

Index 319

Tables

1. System requirements	5	31. Mechanisms used to gather KLZNET attributes	268
2. Capacity planning for historical data logged by component	104	32. Mechanisms used to gather KLZNFS attributes	269
3. Attribute groups supported by the data model	200	33. Mechanisms used to gather KLZPROC attributes	271
4. Information to gather before contacting IBM Software Support	207	34. Mechanisms used to gather KLZPUSR attributes	275
5. Log file management on UNIX compared to log file management on Windows.	208	35. Mechanisms used to gather KLZRPC attributes	276
6. Trace log files for troubleshooting agents	209	36. Mechanisms used to gather KLZSOCKD attributes	276
7. Problems and solutions for installation and configuration	212	37. Mechanisms used to gather KLZSOCKS attributes	277
8. General problems and solutions for uninstallation	215	38. Mechanisms used to gather KLZSWPRT attributes	277
9. Agent problems and solutions	218	39. Mechanisms used to gather KLZSYS attributes	277
10. Tivoli Enterprise Portal problems and solutions	220	40. Mechanisms used to gather KLZVM attributes	278
11. Remote deployment problems and solutions	220	41. Mechanisms used to gather LNXALLUSR attributes	280
12. Tivoli Common Reporting problems and solutions	221	42. Mechanisms used to gather LNXCPU attributes	280
13. Specific situation problems and solutions	221	43. Mechanisms used to gather LNXCPUAVG attributes	280
14. Problems with configuring situations that you solve in the Situation Editor	223	44. Mechanisms used to gather LNXCPUCON attributes	281
15. Problems with configuration of situations that you solve in the Workspace area	224	45. Mechanisms used to gather LNXDISK attributes	282
16. Problems with configuration of situations that you solve in the Manage Tivoli Enterprise Monitoring Services window	225	46. Mechanisms used to gather LNXDSKIO attributes	282
17. Time periods and suffixes for summary tables and views.	227	47. Mechanisms used to gather LNXDU attributes	283
18. Additional columns to report summarization information	228	48. Mechanisms used to gather LNXFILCMP attributes	283
19. Primary key and warehouse changes for the Monitoring Agent for Linux OS	230	49. Mechanisms used to gather LNXFILE attributes	284
20. Scripts for affected attribute groups and summary tables for the Monitoring Agent for Linux OS	232	50. Mechanisms used to gather LNXFILPAT attributes	284
21. Overview of Distributed Monitoring migrated situations	237	51. Mechanisms used to gather LNXGROUP attributes	285
22. Overview of attribute groups to event classes and slots	240	52. Mechanisms used to gather LNXIOEXT attributes	285
23. Mechanisms used to gather KLZCPU attributes	263	53. Mechanisms used to gather LNXIPADDR attributes	286
24. Mechanisms used to gather KLZCPUAVG attributes	264	54. Mechanisms used to gather LNXLOGIN attributes	286
25. Mechanisms used to gather KLZDISK attributes	265	55. Mechanisms used to gather LNXMACHIN attributes	287
26. Mechanisms used to gather KLZDSKIO attributes	265	56. Mechanisms used to gather LNXNET attributes	287
27. Mechanisms used to gather KLZDU attributes	266	57. Mechanisms used to gather LNXNFS attributes	288
28. Mechanisms used to gather KLZIOEXT attributes	266	58. Mechanisms used to gather LNXOSCON attributes	290
29. Mechanisms used to gather KLZLOGIN attributes	267		
30. Mechanisms used to gather KLZLPAR attributes	268		

59. Mechanisms used to gather LNXPING attributes	290	64. Mechanisms used to gather LNXSOCKS attributes	294
60. Mechanisms used to gather LNXPROC attributes	291	65. Mechanisms used to gather LNXSWPRT attributes	295
61. Mechanisms used to gather LNXPUSE attributes	292	66. Mechanisms used to gather LNXSYS attributes	295
62. Mechanisms used to gather LNXRPC attributes	293	67. Mechanisms used to gather LNXVM attributes	296
63. Mechanisms used to gather LNXSOCKD attributes	294		

Chapter 1. Using the monitoring agent

As part of the Tivoli Enterprise Portal for Distributed Systems, the Monitoring Agent for Linux OS offers a central point of management of Linux-based environments. It provides a comprehensive means for gathering exactly the information you need to detect problems early and to prevent them. Information is standardized across all systems, and you can monitor servers from a single workstation. The Tivoli Enterprise Portal lets you easily collect and analyze specific information.

The Monitoring Agent for Linux OS is an intelligent, remote monitoring agent that resides on managed resources. It assists you in anticipating trouble and warns systems administrators when critical events take place on their systems. With the Monitoring Agent for Linux OS, systems administrators can set threshold levels and flags as desired to alert them when the system reaches these thresholds.

For Tivoli Enterprise Portal, information is displayed in named workspaces. Tivoli Enterprise Portal refers to this tabular format for information as a table view. Information can also be displayed in the workspace as charts, graphs, or other formats that you can specify.

The Monitoring Agent for Linux OS provides the following benefits:

- Simplifies application and system management by managing applications, platforms, and resources across your environment.
- Helps to increase 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 new platforms by supporting a wide variety of platforms.
- Improves system performance by letting you integrate, monitor, and manage your system, network, console, and mission-critical applications. A monitoring agent alerts the Tivoli Enterprise Monitoring Server when conditions on the system network meet threshold-based conditions. These alerts notify your systems administrator to limit and control database usage. You can view data gathered by the Tivoli Enterprise Monitoring Server in tables and charts for the status of your distributed database systems.
- Enhances efficiency by monitoring diverse platforms and networks from a single PC screen. Depending on your Tivoli Enterprise Portal configuration, you can collect and monitor data across platforms. Management agents gather and filter status information at the managed resource rather than at the hub, eliminating unnecessary data transmission and sending only data that is relevant to changes in status conditions. The Monitoring Agent for Linux OS helps you monitor and gather the consistent, accurate, and timely information you require to effectively perform your job.

New in this release

For version 6.3 of the monitoring agent, the following enhancements have been made:

- New attributes, Processes Waiting RunTime and Processes Blocked State, added to the VM Stats attribute group. The Virtual Memory Information view of the

Virtual Memory Statistics workspace displays data for the new attributes, Processes Waiting RunTime and Processes Blocked State.

- New attributes, LPAR Name, Maximum Physical CPUs, Number of Physical CPUs, added to the LPAR attribute group.
 - New attribute, Virtual Machine Identifier, added to the Machine Information attribute group.
 - For attribute values calculated as an average of the cumulative CPU ticks between two samples, note that the sample time differs depending on how the agent is invoked to return the values. If the agent is invoked to return the values on-demand (for example, after a workspace refresh), the sample time is 30 secs. If, however, the agent is invoked to return the values by a situation or an historical collection request, the sample time is the same as that of the situation or of the collection interval. The affected attributes include:
 - CPU attribute group: User CPU (Percent), User Nice CPU (Percent), System CPU (Percent), Idle CPU (Percent), Busy CPU (Percent), I/O Wait (Percent), User to System CPU (Percent), and Steal CPU (Percent) attributes
 - CPU Averages attribute group: Estimated Days Until CPU Upgrade, Total CPU Used Current Average (Percent), Total CPU Used Moving Average (Percent), User Nice CPU Current Average (Percent), User Nice CPU Moving Average (Percent), User CPU Current Average (Percent), User CPU Moving Average (Percent), System CPU Current Average (Percent), System CPU Moving Average (Percent), Idle CPU (Percent), Idle CPU Moving Average (Percent), Wait CPU (Percent), and Wait CPU Moving Average (Percent) attributes
 - LPAR attribute group: Total Steal Time (Percent) attribute
 - Process attribute group: Process Instant Busy CPU (Percent) attribute
- You can customize the attribute values by specifying two variables in the lz.ini file: KLZ_CPUSTAT_SAMPLE_SECS for the total CPU metrics (default value: 30 seconds) and KLZ_PROCESS_SAMPLE_SECS for the CPU metrics per process (default value: 60 seconds). If these variables are set to 0, the sampling interval is variable: the samples are taken when the requests come to the agent (for example, at each workspace refresh), and the sampling interval is the difference in time between last two samples (with a minimum of at least 5 seconds).
- The Summarization and Pruning agent automatically creates and maintains the shared dimensions tables. For instructions to enable this feature, see “Configuring the Summarization and Pruning agent to maintain the dimension tables” in the *IBM Tivoli Monitoring Administrator's Guide*. To enhance this feature for the OS Agents Reports package, the installer now prompts you to provide JDBC connection details and credentials for the TDW database. This RegisterPackage script execution step inserts data into the WAREHOUSETCRCONTROL table. After this step, the MANAGEDSYSTEM table and the TIME_DIMENSION table are kept up to date automatically by the Summarization and Pruning agent. However, if you opt not to use this feature and prefer, instead, to manually maintain the dimensions tables, skip this step. For instructions to perform any required manual steps, see “Manually creating and maintaining the dimension tables” in the *IBM Tivoli Monitoring Administrator's Guide*.
 - The agent provides ComputerSystem and IPAddress resources for the Open Services for Lifecycle Collaboration Performance Monitoring (OSLC-PM) service provider. The service provider registers monitoring resources with the Registry Services. Registry Services is a Jazz for Service Management integration service that provides a shared data repository for products in an integrated service management environment.

- The IBM Tivoli Monitoring Infrastructure Management Dashboards for Servers is a web-based application that runs in the Dashboard Application Services Hub. The server dashboards give the overall status of the service areas in your managed network. Use the server dashboards to assess the event and system status of your managed network that is filtered by your area of responsibility. The information ranges from a high-level overview of all managed system groups and the situation events associated with them, to more detailed dashboards with key performance information about the selected group, managed system, or situation event.

Using IBM Tivoli Monitoring

IBM Tivoli Monitoring is the base software for the Monitoring Agent for Linux 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 307 for complete information about IBM Tivoli Monitoring and the Tivoli Enterprise Portal.

Components of the monitoring agent

After you install the Monitoring Agent for Linux OS (product code "klz" or "lz") 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 Linux 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 *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

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

Table 1. System requirements

Operating system versions	System requirements
Linux on zSeries®	<ul style="list-style-type: none"> • RedHat Enterprise Linux 5 on zSeries (64 bit) • RedHat Enterprise Linux 5 on zSeries (31 bit) • RedHat Enterprise Linux 6 on zSeries (31 bit) • RedHat Enterprise Linux 6 on zSeries (64 bit) • SuSE Linux Enterprise Server 10 for zSeries (64 bit) • SuSE Linux Enterprise Server 11 for zSeries (64 bit)
Linux on Intel (32-bit)	<ul style="list-style-type: none"> • Asianux 3.0 for Intel x86-32 (32 bit) • Red Flag 5.0 for Intel x86-32 (32 bit) • RedHat Enterprise Linux 5 Intel x86-32 (32 bit) • RedHat Enterprise Linux 6 Intel x86-32 (32 bit) • SuSE Linux Enterprise Server 10 Intel x86-32 (32 bit) • SuSE Linux Enterprise Server 11 on Intel x86-32 (32 bit)
Linux on iSeries® and pSeries	<ul style="list-style-type: none"> • RedHat Enterprise Linux 5 on iSeries and pSeries (64 bit) • RedHat Enterprise Linux 6 on iSeries and pSeries (64 bit) • SuSE Linux Enterprise Server 10 for iSeries and pSeries (64 bit) • SuSE Linux Enterprise Server 11 for iSeries and pSeries (64 bit)
Linux on x86-64	<ul style="list-style-type: none"> • Asianux 3.0 on Intel x86-64 (64 bit) • RedHat Enterprise Linux 5 on Intel x86-64 (64 bit) • RedHat Enterprise Linux 6 on Intel x86-64 (64 bit) • SuSE Linux Enterprise Server 10 on Intel x86-64 (64 bit) • SuSE Linux Enterprise Server 11 on Intel x86-64 (64 bit)
Natively on VMWare	<ul style="list-style-type: none"> • ESX Server 3.0.1 x86-32 • ESX Server 3.0.1 x86-64 • ESX Server 3.5 x86-32 • ESX Server 3.5 x86-64 • ESX Server 4.0 x86-32 • ESX Server 4.0 x86-64
The Linux version must support the Korn shell (ksh) and Motif Window Manager (libmotif) for installation of the monitoring agent.	
Memory	<ul style="list-style-type: none"> • 30 MB RAM for the Monitoring Agent for Linux OS

Table 1. System requirements (continued)

Operating system versions	System requirements
<p>Disk space</p>	<p>The Monitoring Agent for LINUX OS needs 180 MB of disk space in the file system where it is to be installed through the local install method. It needs 135 MB of disk space in the /tmp file system and 185 MB of disk space in the file system where the agent is to be installed through the tacmd createNode command. It needs 268 MB of disk space when it is updated by using the command tacmd updateAgent.</p> <p>For historical data disk space information, see “Disk capacity planning for historical data” on page 103.</p>
<p>Other requirements</p>	<ul style="list-style-type: none"> • 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. • 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. • Linux versions require some compatibility libraries to be installed for the agent to work correctly. The latest versions of libstdc++, libgcc, and compat-libstdc++, are required for the agent to work correctly. ² • On 64-bit systems, both the 32-bit RPM files (or the 31-bit RPM files, in the case of zLinux systems) and the 64-bit RPM files need to be installed.
<p>Note:</p> <p>1. In native 64-bit mode, not tolerance mode.</p>	

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.

The Linux OS Monitoring Agent requires the installation of the latest versions of the following libraries:

- libstdc++
- libgcc
- compat-libstdc++

These libraries are available on the Linux operating system installation media and Service Packs. Each library can have multiple packages and each should be installed. If you are on a 64-bit system you must have the 32-bit and 64-bit versions of these libraries.

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

Naming instances

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
- Machine 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 by 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.

See "Unique names for monitoring components" on page 216 for more information about creating unique names.

Running as a non-Administrator user

The Monitoring Agent for Linux OS can be run by a non-Administrator user (a non-root user), however some functionality becomes unavailable. The Machine BIOS information uses the `dmidecode` executable to extract the relevant information. This Linux provided executable must be run by the Administrator user to extract BIOS information. This attribute group does not report data if the agent is not run by the Administrator user. This information is also used by Tivoli Application Dependency Discovery Manager.

A non-Administrator user can only access the directories that it has permissions to read. Therefore, functionality of the File Information attribute group might be reduced.

For 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 non-Administrator user is monitoring agents running as a different non-Administrator user.

Also for Agent Management Services, the watchdog cannot stop or start any agent that it does not have privileges to stop or start. If the OS agent is running as a user other than Administrator but you would still like to use it to stop and start other agents, the sudo facility on UNIX and Linux provides one way of supporting this capability. In the example that follows, the OS agent user is a member of a group called 'itm'. Also, it is assumed that Agent Management Services will not be prompted for a password to perform these operations and that the target agents' user IDs are 'user1' and 'user2':

```
# sudoers file.
#
# This file MUST be edited with the 'visudo' command as root.
# Failure to use 'visudo' may result in syntax or file permission errors
# that prevent sudo from running.
#
# See the sudoers man page for the details on how to write a sudoers file.
#

# Host alias specification

# User alias specification

# Cmnd alias specification
Cmnd_Alias AMSAGENTSTART = /opt/PAS/ITMTEST/bin/itmcmd agent -[po] [[\:alnum\:]_]*
start [[\:alnum\:]] [[\:alnum\:]],/opt/PAS/ITMTEST/bin/itmcmd agent start
[[\:alnum\:]] [[\:alnum\:]]
Cmnd_Alias AMSAGENTSTOP = /opt/PAS/ITMTEST/bin/itmcmd agent -[po] [[\:alnum\:]_]*
stop [[\:alnum\:]] [[\:alnum\:]],/opt/PAS/ITMTEST/bin/itmcmd agent stop
[[\:alnum\:]] [[\:alnum\:]]
Cmnd_Alias ITMAMSCMD = AMSAGENTSTART,AMSAGENTSTOP
# Defaults specification

# Runas alias specification

Runas_Alias ITMAGENTIDS = user1,user2

# Same thing without a password
%itmusers ALL=( ITMAGENTIDS ) NOPASSWD: ITMAMSCMD
```

This is just one possible example. The sudo facility has many advanced capabilities including the ability to audit and to alert administrators of usage of the sudo command by unauthorized users. See your operating system's sudo man pages for more information.

In the agentInstanceCommand.sh script, replace calls to 'su' with calls to 'sudo'. For example:

```
if [ -z "$USR" ]; then
    $START_CMD
else
    # su - $USR -c "$START_CMD"
    sudo -u $USR $START_CMD
fi
```

```

...
if [ -z "$USR" ]; then
    $STOP_CMD
else
    # su - $USR -c "$STOP_CMD"
    sudo -u $USR $STOP_CMD
fi

```

Ensure that the user1 and user2 users also have write permission to any files to which an application agent needs to write.

Using Agent Management Services

There are two watchdog monitors that run as part of the Monitoring Agent for Linux. 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 watches the OS Agent, so as long as its Availability Status is showing 'Running' in the Agents' 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, it is able to respond to Tivoli Enterprise Portal Desktop queries and Take Actions performed against these other agent processes. This is the data that is seen 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 running as part of the OS Agent. These are non-OS agents, so the Monitoring Agent for Linux 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 Linux is started. This watchdog does not have a communication facility, so it cannot report information to the Tivoli Enterprise Portal or respond to Take Actions. It is not an agent per se, 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 lz disarmWatchdog.sh command. This 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 lz 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 lz 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 has been added to the Process attribute group. Its content, a regular expression, is sent to the agent as a filter object and is intended to act only on the Process Command (Unicode) attribute. For example, the agent uses the value provided in the Process Filter attribute to match with 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.

To use this enhancement, create queries and situations on the 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

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 or delete the predefined workspaces, but you can create new workspaces by editing them and saving the changes with a different name.

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 and the information in that section for each individual workspace.

Organization of the predefined workspaces

The following list shows the organization of the predefined workspaces provided with IBM Tivoli Monitoring: Linux OS Agent.

- “Capacity Usage Information workspace” on page 13
 - “CPU Averages workspace” on page 14
 - “Virtual Memory Usage Trends workspace” on page 24
- “Disk Usage workspace” on page 15
- “File Information workspace” on page 15
 - “All Files workspace” on page 13
- “Network workspace” on page 19
 - “Sockets Information workspace” on page 21
 - “NFS Statistics workspace” on page 20
 - “RPC Statistics workspace” on page 21
- “Process workspace” on page 20
 - “Process CPU Usage workspace” on page 20
 - “Process User Information workspace” on page 21
- “System Information workspace” on page 22
 - “System Configuration workspace” on page 22
 - “Disk I/O Rate workspace” on page 14
 - “Disk I/O Extended Rate workspace” on page 14
 - “Virtual Memory Statistics workspace” on page 23
- “Users Workspace” on page 23
- “Agent Management Services workspace” on page 12
 - “Agents' Management Log workspace” on page 12

This agent also includes the following linked workspaces:

- Historical Summarized Availability
- Historical Summarized Availability Daily
- Historical Summarized Availability Hourly
- Historical Summarized Availability Weekly
- Historical Summarized Capacity
- Historical Summarized Capacity Daily
- Historical Summarized Capacity Hourly
- Historical Summarized Capacity Weekly
- Historical Summarized Performance
- Historical Summarized Performance Daily
- Historical Summarized Performance Hourly
- Historical Summarized Performance Weekly

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.

The remaining sections contain descriptions of each of these predefined workspaces, which are organized alphabetically within the group.

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 Linux.

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.

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 `lz.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.

All Files workspace

The All Files workspace is reached by right-clicking the File Information navigator item in the Tivoli Enterprise Portal. The views are:

- File Size - Top Ten (bar chart)
- All Files (table view)

The File Size - Top Ten bar chart displays the sizes of the largest files. The All Files table provides file information.

Capacity Usage Information workspace

The Capacity Usage Information workspace reflects the “health” of your system by providing CPU, disk, and swap space usage statistics. 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 is comprised of three views. The views are:

- Disk Usage Averages (table view)
- Disk Space Usage (bar chart)
- Disk Usage Averages (bar chart)

The Disk Usage Averages table provides information on the system's current disk usage. The Disk Space Usage bar chart displays the system's current disk usage. The Disk Usage Averages bar chart displays average disk usage information. With

the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

CPU Averages workspace

The CPU Averages workspace is reached by right-clicking the Capacity Usage Information navigator item in the Tivoli Enterprise Portal. 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). The workspace is comprised of 3 views. The views are Current Overall CPU Usage bar chart, CPU Averages (Hourly Updates) chart, and CPU Usage Trends table.

Disk I/O Extended Rate workspace

The Disk I/O Extended Rate workspace is reached by right-clicking the System Information navigator item in the Tivoli Enterprise Portal. The Disk I/O Extended Rate workspace provides detailed input/output statistics and "calculations", including the queue length and size in sectors of read and write requests, the rate of those requests, and wait times associated with requests. 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 is comprised of two views. The views are:

- Disk I/O Extended Rate (table view)
- Disk Service Time (bar chart)
- Disk Activity (bar chart)

The Disk I/O Extended Rate table details the input/out data and calculated values associated with disk activity. The Disk Service Time chart displays average services time in minutes. The Disk Activity chart displays read and write sectors in seconds. Based on the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

Note: The attributes associated with this workspace are only available for systems with a 2.4 (or higher) kernel.

Disk I/O Rate workspace

The Disk I/O Rate workspace is reached by right-clicking the System Information navigator item in the Tivoli Enterprise Portal. The Disk I/O Rate workspace provides input/output statistics, including the transfer rates, block read rates, and block write rates of your monitored systems. 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 is comprised of two views. The views are:

- Disk I/O Rate (table view)
- Disk I/O Rate (bar chart)

The Disk I/O Rate table includes transfer rates, block read rates, and block write rates for your monitored systems. The Disk I/O Rate chart provides "at a glance" rate details associated with disk reads, writes, and transfers. Based on the

information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

Note: The attributes associated with this workspace are only available for systems with a 2.4 (or higher) kernel.

Disk Usage workspace

The Disk Usage workspace reflects the health of storage space within your monitored systems. 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 is comprised of four views. The views are:

- Space Used Percent (bar chart)
- Inodes Used Percent (bar chart)
- Disk Space (bar chart)
- Disk Usage (table view)

The three charts that comprise this workspace provide “at a glance” percentages of the space used, percentages of the inodes used, and amounts of disk space used/available for each monitored disk. The Disk Usage table captures this information, as well as mount point and file system data, in tabular form. Based on the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

File Information workspace

The File Information workspace displays file information statistics. The views are:

- File Size - Top Ten (bar chart)
- File Size - Top Ten (table view)

Historical Summarized Availability workspace

The Historical Summarized Availability workspace shows the percentage of time that a managed resource was available during the number of months that you specify in the Time Span dialog. This workspace consists of the following two graphical views:

- Availability (average over months), which shows the percentage of time that managed resources were available, grouped by server
- Process Summary (average over months), which shows the percentage of time per system that each process was used by the server

Historical Summarized Availability Daily workspace

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

- Availability (daily), which shows the percentage of the day that the server was available, summarized by day
- Process Summary (daily), which shows details such as memory and processor usage for processes that were running on the server, summarized by day

- System Summary (daily), which shows system information for the server, such as the operating system type, name, version, and manufacturer, summarized by day

Historical Summarized Availability Hourly workspace

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

- Availability (hourly), which shows the percentage of the hour that the server was available, summarized by hour
- Process Summary (hourly), which shows details such as memory and processor usage for processes that were running on the server, summarized by hour
- System Summary (hourly), which shows system information for the server, such as the operating system type, name, version, and manufacturer, summarized by hour

Historical Summarized Availability Weekly workspace

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

- Availability (weekly), which shows the percentage of system time that the server was available, summarized by week
- Process Summary (weekly), which shows processes that kept the server busy, summarized by week
- System Summary (weekly), which shows system information such as the operating system type, name, version, and manufacturer, summarized by week

Historical Summarized Capacity workspace

The Historical Summarized Capacity workspace shows usage of system resources during the time span that you specify in the Time Span dialog. This workspace consists of the following 5 graphical views:

- Network Interface Activity (average over months), which shows network traffic for the server for all network interfaces on the system during the time span that you specify in the Time Span dialog
- Processor Utilization (average over months), which shows CPU usage, including idle CPU time, for all processors that are associated with the server during the specified time period
- Memory Utilization (average over months), which shows memory used, free memory, and swapped memory use during the specified time period
- Disk Utilization (maximum over months), which shows the maximum percentage of space used on the system's logical disks during the specified time period
- Disk Capacity (minimum over months), which shows information about the remaining number of days until the disk is full based on the current rate of disk usage, and the remaining number of days until the disk is full based on peak rate of disk usage, for all disks that are associated with the server

Historical Summarized Capacity Daily workspace

The Historical Summarized Capacity Daily workspace shows system usage summarized by day. This workspace consists of the following four graphical views:

- Network Interface Activity, which shows network traffic for the server, including packet collision rates, during the specified time period, summarized by day
- Processor Utilization, which shows CPU usage (including an idle, busy, or waiting CPU), for all processors that are associated with the server during the specified time period, summarized by day
- Memory Utilization, which shows memory used, free memory, and swapped memory use during the specified time period, summarized by day
- Disk Utilization, which shows percentage of space used or available on the system's logical disks during the specified time period, summarized by day

Historical Summarized Capacity Hourly workspace

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

- Network Interface Activity, which shows network traffic, including collisions, packet transmittal and count transmittal for the server during the specified time period, summarized by hour
- Processor Utilization, which shows average CPU usage (idle, busy, and waiting), for all processors that are associated with the server during the specified time period, summarized by hour
- Memory Utilization, which shows memory used, free memory, and swapped memory use during the specified time period, summarized by hour
- Disk Utilization, which shows percentages of space used and available on all the system's logical disks during the specified time period, summarized by hour

Historical Summarized Capacity Weekly workspace

The Historical Summarized Capacity Weekly workspace shows system resources used, summarized by week. This workspace consists of the following 5 graphical views:

- Network Interface Activity, which shows network traffic for the server during the specified time period, summarized by week
- Processor Utilization, which shows CPU usage, especially idle CPU time, for all processors that are associated with the server during the specified time period, summarized by week
- Maximum Memory Utilization, which shows maximum memory used, free memory, and swapped memory during the specified time period, summarized by week
- Average Memory Utilization, which shows average memory that the server used during the specified time period, summarized by week
- Disk Utilization, which shows the maximum percentage of space used on all the system's logical disks during the specified time period, summarized by week

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 5 graphical views:

- Network Activity (average over months), which shows (in the sample period) percentages of errors and collisions in network traffic for all networks that are associated with the system during the time span that you specify in the Time Span dialog
- System Load (average over months), which shows the system workload during the specified time period
- Disk I/O Traffic (average over months), which shows the average percentage of time that the disk was busy during the specified time period
- Memory Page Faults (average over months), which shows the average rate of page in and page out for the system during the specified time period
- Processor Performance (average over months), which shows the average percentage of usage that users consumed and the average processor waiting time for the server during the specified time period

Historical Summarized Performance Daily workspace

The Historical Summarized Performance Daily workspace shows the performance of system resources, summarized by day. This workspace consists of the following 5 graphical views:

- Network Activity (daily), which shows the average network activity for a server, including transmittals, packet collisions, carrier losses, and so on, summarized by day
- System Load (daily), which shows the system workload during the specified time period, summarized by day
- Disk I/O Traffic (daily), which shows the average percentage of time that the disk was busy during the specified time period, summarized by day
- Memory Page Faults (daily), which shows the average rate of page in and page out for the system during the specified time period, summarized by day
- Processor Performance (daily), which shows the percentage of processor time that users consumed, as well as the waiting time that the CPU spent during the specified time period, summarized by day

Historical Summarized Performance Hourly workspace

The Historical Summarized Performance Hourly workspace shows the performance of system resources, summarized by hour. This workspace consists of the following 5 graphical views:

- Network Activity (hourly), which shows the network activity for a server, including transmittals, packet collisions, carrier losses, and so on, summarized by hour
- System Load (hourly), which shows the system workload during the specified time period, summarized by hour
- Disk I/O Traffic (hourly), which shows the average percentage of time that the disk was busy during the specified time period, summarized by hour
- Memory Page Faults (hourly), which shows the average rate of page in and page out for the system during the specified time period, summarized by hour

- Processor Performance (hourly), which shows the percentage of processor time that users consumed, as well as the waiting time that the CPU spent during the specified time period, summarized by hour

Historical Summarized Performance Weekly workspace

The Historical Summarized Performance Weekly workspace shows the performance of system resources, summarized by week. This workspace consists of the following 5 graphical views:

- Network Activity (weekly), which shows the network activity for a server, including errors and packet collisions, for all networks associated with the server, summarized by week
- System Load (weekly), which shows the system workload during the specified time period, summarized by week
- Memory Page Faults (weekly), which shows the average rate of page in and page out for the system during the specified time period, summarized by week
- Disk I/O Traffic (weekly), which shows the average percentage of time that the disk was busy during the specified time period, summarized by week
- Processor Performance (weekly), which shows the percentage of processor time that users consumed, as well as the waiting time that the CPU spent during the specified time period, summarized by week

Linux workspace

The Linux workspace reflects the health of the system. 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 is comprised of three views. The views are:

- System CPU Usage (bar chart)
- Disk IO Transfers (bar chart)
- System Load Averages (bar chart)

Network workspace

The Network workspace reflects the health of the network components within your monitored systems. 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 is comprised of four views. The views are:

- Network Errors (bar chart)
- Network Activity (bar chart)
- Network Devices (table view)
- IP Addresses (table view)

The Network Errors chart shows the number of input errors, output errors, and collisions for the sampling period. The Network Activity chart shows the number of packets received and transmitted per second. The Network Devices table reflects your network's performance based on its transmission, reception, and collision data. The IP Addresses table shows the IP addresses of the network interface

names. Based on the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

NFS Statistics workspace

The NFS Statistics workspace is reached by right-clicking the Network navigator item in the Tivoli Enterprise Portal. The NFS Statistics workspace provides statistics on the operations involving the Network File System, such as the number and type of calls being made, and the percentages those types of calls make up in relation to total calls. 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). The views are:

- Network Errors (bar chart)
- RPC Network Activity (bar chart)
- NFS Statistics (table view)

Based on the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

Process workspace

The Process workspace reflects the health of specific processes within your monitored systems. 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 is comprised of three views. The views are:

- Process CPU Percent Usage (bar chart)
- Process + Child CPU Percent Usage (bar chart)
- Process Information Detail (table view)

The Process CPU Percent Usage chart displays the percentage of CPU time spent in kernel mode and spent in user mode by process. The Process + Child CPU Percent Usage chart displays the cumulative percentage of CPU time spent in kernel mode and spent in user mode. The Process Information Detail table lists in tabular form a wide range of process characteristics such as data set size, kernel scheduling priority, the number of pages of memory, and the number of page faults. Based on the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

Process CPU Usage workspace

The Process CPU Usage workspace reflects the CPU consumption of specific processes within your monitored systems. Use this workspace to check systems performances that could be at risk in case of processes consuming too much CPU. Existing percentages of CPU usage are computed on the lifetime of the processes and are unable to report sudden peaks. The views are:

- Process Instant CPU Percent Usage
- Process CPU Usage Detail

Process User Information workspace

The Process User Information workspace is reached by right-clicking the Process navigator item in the Tivoli Enterprise Portal. The Process User Information workspace identifies process owners of your monitored Linux system and details their usage. 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 is comprised of three views. The views are:

- Process CPU Percent Usage (bar chart)
- Process + Child CPU Percent Usage (bar chart)
- Process User Information (table view)

The Process CPU Percent Usage chart displays the percentage of CPU time spent in kernel mode and spent in user mode by process. The Process + Child CPU Percent Usage chart displays the cumulative percentage of CPU time spent in kernel mode and spent in user mode. The Process User Information table provides in tabular form the names of effective groups, file system groups, real groups, and saved groups for your monitored systems. Based on the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

RPC Statistics workspace

The RPC Statistics workspace is reached by right-clicking the Network navigator item in the Tivoli Enterprise Portal. The RPC (remote procedure call) workspace provides statistics on the number and type of calls being made to the server and clients, including statistics on the number of calls that are not valid or had to be retransmitted. 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). The views are:

- Network Errors (bar chart)
- RPC Network Activity (bar chart)
- RPC Statistics (table view)

Based on the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

Sockets Information workspace

The Sockets Information workspace is reached by right-clicking the Network navigator item in the Tivoli Enterprise Portal. The Sockets Information workspace reflects the health of the socket connections within your monitored systems. 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 is comprised of three views. The views are:

- Sockets Used by Protocol (bar chart)
- Network Activity (bar chart)

- Socket Services Information (table view)

The Sockets Used by Protocol chart shows a count of the sockets currently in use and the high water mark for each protocol during the sampling period. The Network Activity chart shows the number of packets received and transmitted per second. The Socket Services Information table provides a detailed perspective of each socket that you are monitoring. Based on the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

Specific File Information workspace

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 contains detailed information about a specific file or directory. You can access this information down through the lowest directory structure. This workspace is comprised of two views. The views are:

- File Information (table view)
- Take Action view

System Configuration workspace

The System Configuration workspace is reached by right-clicking the System Information workspace in the Tivoli Enterprise Portal. The System Configuration workspace displays information about CPU usage, the processor's configuration, and operating system level. 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). It contains the following views:

- CPU Usage (bar chart)
- Processor Configuration Information (table view)
- OS Version Information (table view)
- LPAR Information (table view)

System Information workspace

The System Information workspace reflects the health of your monitored systems by displaying data associated with system loads, context switching, and process creation. 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 is comprised of the following views:

- CPU Usage (bar chart)
- Paging Rates (bar chart)
- System Load (bar chart)
- Virtual Memory Statistics (bar chart)
- System Statistics (table view)

The CPU Usage bar chart shows the percentage of idle CPU time, system CPU time, user CPU time, user nice CPU time, and steal CPU time of the monitored processor. The System Load chart depicts the load on your monitored system's processor during the previous one, five, and 15 minutes. The paging rates chart displays information about paging in and out as well as swapping in and out

trends in seconds. The Virtual Memory Statistics chart depicts the current usage and availability of a variety of memory categories (buffered, cached, shared, and swapped). The System Statistics table lists in tabular form the source data of these charts and gauge. Based on the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

The System Configuration workspace is reached by right-clicking the System Information navigator item in the Tivoli Enterprise Portal.

Users Workspace

The Users workspace identifies logged in users. 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 is comprised of three views. The views are:

- Process User Information (table view)
- Total User Logins (needle gauge)
- User Login Information (table view)

The Process User Information table provides in tabular form the names of effective groups, file system groups, real groups, and saved groups for your monitored systems. The Total User Logins gauge displays the number of users logged on to the monitored system during the monitoring period. The User Login Information table lists users, their log in time, and their idle time. Based on the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

Virtual Memory Statistics workspace

The Virtual Memory Statistics workspace is reached by right-clicking the System Information navigator item in the Tivoli Enterprise Portal. The Virtual Memory Statistics workspace provides a snapshot of your monitored systems memory usage. 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 is comprised of four views. The views are:

- Context Switches Percent Change (needle gauge)
- Context Switches Per Second (needle gauge)
- Virtual Memory Statistics (bar chart)
- Virtual Memory Information (table view)

The Context Switches Percent Change gauge reflects the percentage change in the number of context switches per second. The Context Switches Per Second gauge shows the number of context switches per second. The Virtual Memory Statistics chart depicts the current usage and availability of a variety of memory categories (buffered, cached, shared, and swapped). The Virtual Memory Information table presents the Virtual Memory Usage information in tabular form. Based on the information that this workspace provides, you can recommend changes, set up situations, and verify that your recommended changes improve performance.

Virtual Memory Usage Trends workspace

The Virtual Memory Usage Trends workspace is reached by right-clicking the Capacity Usage Information navigator item in the Tivoli Enterprise Portal. 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). The views are:

- Current Virtual Memory Usage (bar chart)
- Virtual Memory Averages (bar chart)
- Swap Space Usage Trends (table view)

The Current Virtual Memory Usage bar chart displays memory usage information. The Virtual Memory Averages bar chart displays virtual memory usage trend information. The Swap Space Usage Trends table provides several types of swap space information.

Chapter 4. Attributes

Attributes are the application properties that are measured and reported by the monitoring agent, 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 and registers an *event* if the condition is met. You are alerted to events by indicator icons that are displayed in the Navigator.

Note: 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.

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

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 monitoring agent attributes to build situations that monitor the performance of your 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.

Attribute groups

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 the monitoring agent. The groups are collected in attribute tables that are designated in brackets [] after the group name.

- Agent Availability Management Status [KLZPASMGMGT]
- Agent Active Runtime Status [KLZPASSTAT]
- Alerts Table [KLZPASALRT]
- All Users Group [LNXALLUSR]
- Configuration Information [KLZPASCAP]
- CPU [KLZCPU]
- CPU (superseded) [LNXCPU]
- CPU Averages [KLZCPUAVG]
- CPU Averages (superseded) [LNXCPUAVG]
- CPU Configuration [LNXCPUCON]
- Disk [KLZDISK]
- Disk (superseded) [LNXDISK]
- Disk IO [KLZDSKIO]
- Disk IO (superseded) [LNXDSKIO]
- Disk Usage Trends [KLZDU]
- Disk Usage Trends (superseded) [LNXDU]
- File Comparison Group [LNXFILCMP]
- File Information [LNXFILE]
- File Pattern Group [LNXFILPAT]
- I/O Ext [KLZIOEXT]
- I/O Ext (superseded) [LNXIOEXT]
- IP Address [LNXIPADDR]
- Linux Group [LNXGROUP]
- Linux Host Availability [LNXPING]

- Linux TCP Statistics [KLZTCP]
- LPAR [KLZLPAR]
- Machine Information [LNXMACHIN]
- Network [KLZNET]
- Network (superseded) [LNXNET]
- NFS Statistics [KLZNFS]
- NFS Statistics (superseded) [LNXNFS]
- OS Configuration [LNXOSCON]
- Process [KLZPROC]
- Process (superseded) [LNXPROC]
- Process User Info [KLZPUSR]
- Process User Info (superseded) [LNXPUSR]
- RPC Statistics [KLZRPC]
- RPC Statistics (superseded) [LNXRPC]
- Sockets Detail [KLZSOCKD]
- Sockets Detail (superseded) [LNXSOCKD]
- Sockets Status [KLZSOCKS]
- Sockets Status (superseded) [LNXSOCKS]
- Swap Rate [KLZSWPRT]
- Swap Rate (superseded) [LNXSWPRT]
- System Statistics [KLZSYS]
- System Statistics (superseded) [LNXSYS]
- User Login [KLZLOGIN]
- User Login (superseded) [LNXLOGIN]
- VM Stats [KLZVM]
- VM Stats (superseded) [LNXVM]

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.

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 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 VRM information for the agent.

Build Date The build date information for the agent. Superseded by the Build Number attribute.

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.

Operating System The operating system identification. The following values are valid: Win2000 (0), Win2003 (1), Win2008 (2), AIX (3), Linux (4), UNKNOWN (5), NA (-1).

Server Name The origin node of the collecting agent.

Service Name The service name.

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 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).

Page Faults Per Second The total page faults.

Parent Process ID The parent process ID.

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 process resident size.

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 running managed agent.

Alerts Table attributes

Use the Alerts Table attributes to view exceptional Warning and Critical level events surfaced by Agent Management Services. These events have to do with the operation of Agent Management Services or conditions affecting its ability to manage agents. They include the following:

- Agent stopped abnormally.
- Agent restart failed.
- Agent exceeded restart tries.
- Agent not found.
- Agent exceeded policy defined memory threshold.
- Agent exceeded policy defined CPU threshold.
- Agent manual stop failed.
- Agent removed from system - CAP file removed.

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).

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.

All Users attributes

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

Duplicate User Name True if the user name is listed more than once in /etc/passwd. The following values are valid: 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 (-1), False (0), and True (1).

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

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

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.

Time Stamp 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. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

User ID (Superseded) The numeric ID the system assigned to a user. For example, the following value is valid: Value_Exceeds_Maximum=2147483647.

User Sessions The number of log in sessions this user currently has established. For example, the following value is valid: Value_Exceeds_Maximum=2147483647.

Configuration Information attributes

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

Agent Name The sub agent name.

Agent Path The fully qualified path to 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 maximum CPU allowed.

Dependencies The dependent agents.

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 maximum memory allowed.

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 identification. The following values are valid: Unknown (0), Windows (1), Linux (2).

Operating System Name The operating system name.

Operating System Version The operating system version.

PAS_ID The PAS sub agent ID.

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 service name.

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.

CPU attributes

The CPU attributes refer to processor characteristics such as idle time, system CPU time, and user CPU time.

Busy CPU (Percent) The percentage of time the CPU was busy. The following values are valid: integers in the range 0 to 100. Note that the attribute value is averaged in synch with the situation or historical collection interval.

CPU ID The processor ID. The following values are valid: integers in the range 0 to 999. Use this attribute to determine the processor ID. In a SMP system with more than one processor, the CPU report will show CPU ID as "aggregate" on the first row. This means the data row return aggregated CPU statistics. For example, the following value is valid: Aggregate=-1.

Idle CPU (Percent) Percentage of idle CPU time during the sampling period. The following values are valid: integers in the range 0 to 100. 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. Note that the attribute value is averaged in synch with the situation or historical collection interval.

I/O Wait (Percent) The percentage of time the CPU was in a wait input/output state. The following values are valid: integers in the range of 0 to 100. Note that the attribute value is averaged in synch with the situation or historical collection interval.

Steal CPU (Percent) The percentage of time spent in involuntary wait by the virtual CPU while the hypervisor is servicing another virtual processor (applies only to Linux on PPC P5 or greater). Note: -1 indicates Not Available and -2 indicates Not Collected. Note that the attribute value is averaged in synch with the situation or historical collection interval.

System CPU (Percent) Percentage of system CPU time during the sampling period. The following values are valid: integers in the range 0 to 100. Use this attribute to determine the percentage of system or per processor CPU time devoted to executing Linux system kernel code. System CPU time includes time spent executing system calls and performing administrative functions. Note that the attribute value is averaged in synch with the situation or historical collection interval.

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

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

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.

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

User CPU (Percent) Percentage of user CPU time during the sampling period. The following values are valid: integers in the range 0 to 100. 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. Note that the attribute value is averaged in synch with the situation or historical collection interval.

User Nice CPU (Percent) Percentage of user nice CPU time during the sampling period. The following values are valid: integers in the range 0 to 100. Note that the attribute value is averaged in synch with the situation or historical collection interval.

User to System CPU (Percent) Of the total CPU time, the percentage consumed by users. The following values are valid: -10000 to + 10000. Note that the attribute value is averaged in synch with the situation or historical collection interval.

CPU attributes (superseded)

The CPU attributes refer to processor characteristics such as idle time, system CPU time, and user CPU time. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Busy CPU (Percent) The percentage of time the CPU was busy. The following values are valid: integers. The following values are valid: integers in the range 0 to 100. (Superseded.)

CPU ID The processor ID. The following values are valid: integers in the range 0 to 999. Use this attribute to determine the processor ID. In a SMP system with more than one processor, the CPU report will show CPU ID as "aggregate" on the first row. This means the data row return aggregated CPU statistics. For example, the following value is valid: Aggregate=-1. (Superseded.)

Idle CPU (Percent) Percentage of idle CPU time during the sampling period. The following values are valid: integers in the range 0 to 100. 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. (Superseded.)

I/O Wait (Percent) The percentage of time the CPU was in a wait input/output state. The following values are valid: integers in the range of 0 to 100. (Superseded.)

System CPU (Percent) Percentage of system CPU time during the sampling period. The following values are valid: integers in the range 0 to 100. Use this attribute to determine the percentage of system or per processor CPU time devoted to

executing Linux system kernel code. System CPU time includes time spent executing system calls and performing administrative functions. (Superseded.)

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

User CPU (Percent) Percentage of user CPU time during the sampling period. The following values are valid: integers in the range 0 to 100. 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. (Superseded.)

User Nice CPU (Percent) Percentage of user nice CPU time during the sampling period. The following values are valid: integers in the range 0 to 100. (Superseded.)

User to System CPU (Percent) Of the total CPU time, the percentage consumed by users. The following values are valid: -10000 to + 10000. (Superseded.)

CPU Averages attributes

The CPU Averages attributes refer to CPU usage, System CPU time, idle CPU time, user CPU time, and user nice CPU time characteristics.

Estimated Days Until CPU Upgrade The number of days until CPU Usage Moving average hits 100% Rate. The following values are valid: integers. Note: -1 indicates Not Available and -2 indicates Not Collected. Note that the attribute value is averaged in synch with the situation or historical collection interval.

Idle CPU Moving Average (Percent) The moving average of the idle CPU time for the system, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

Idle CPU (Percent) The current average of the idle CPU time for the system, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

System CPU Current Average (Percent) The current average of the System CPU time, expressed as a percentage. The following values are valid: integers between 0

and 100, such as 85 for 85%. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

System CPU Moving Average (Percent) The moving average of the System CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

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

Total CPU Used Current Average (Percent) The current average of CPU usage, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

Total CPU Used Moving Average (Percent) The moving average of CPU usage, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

User CPU Current Average (Percent) The current average of the user CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

User CPU Moving Average (Percent) The moving average of the user CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

User Nice CPU Current Average (Percent) The current average of the user nice CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

User Nice CPU Moving Average (Percent) The moving average of the user nice CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

Wait CPU Moving Average (Percent) The moving current average of the wait CPU time, expressed as a percentage. The following values are valid: integers between 0

and 100. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

Wait CPU (Percent) The current average of the wait CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100. This average is calculated hourly. Note that the attribute value is averaged in synch with the situation or historical collection interval.

CPU Averages attributes (superseded)

The CPU Averages attributes refer to CPU usage, System CPU time, idle CPU time, user CPU time, and user nice CPU time characteristics. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Estimated Days Until CPU Upgrade The number of days until CPU Usage Moving average hits 100% Rate. The following values are valid: integers. Note: -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

Idle CPU Moving Average (Percent) The moving average of the idle CPU time for the system, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. (Superseded.)

Idle CPU (Percent) The current average of the idle CPU time for the system, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. (Superseded.)

System CPU Current Average (Percent) The current average of the System CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. (Superseded.)

System CPU Moving Average (Percent) The moving average of the System CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. (Superseded.)

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*. (Superseded.)

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Total CPU Used Current Average (Percent) The current average of CPU usage, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. (Superseded.)

Total CPU Used Moving Average (Percent) The moving average of CPU usage, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. (Superseded.)

User CPU Current Average (Percent) The current average of the user CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. (Superseded.)

User CPU Moving Average (Percent) The moving average of the user CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. (Superseded.)

User Nice CPU Current Average (Percent) The current average of the user nice CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. (Superseded.)

User Nice CPU Moving Average (Percent) The moving average of the user nice CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly. (Superseded.)

Wait CPU Moving Average (Percent) The moving current average of the wait CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100. This average is calculated hourly. (Superseded.)

Wait CPU (Percent) The current average of the wait CPU time, expressed as a percentage. The following values are valid: integers between 0 and 100. This average is calculated hourly. (Superseded.)

CPU Configuration attributes

The CPU Configuration attributes refer to configuration characteristics such as CPU ID, CPU Family, and Clock Speed.

Model Name The process model name.

Processor Cache Size (KB) The processor cache size (Kb). Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum.

Processor Clock Speed (MHz) The processor clock speed (MHz). Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum.

Processor Family Number The process family number. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Processor ID The processor ID.

Processor Model Number The processor model number. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum.

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

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

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.

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

Processor Vendor ID The Processor Vendor ID.

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.

Disk Free (MB) The amount of free space on a disk, expressed in megabytes. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Disk Free Percent The amount of free space on a disk, expressed as a percentage. Note: the value -1 indicates Not Available and the value -2 indicates Not Collected.

Disk Name The name of the physical disk partition where the file system is mounted. This is the physical location of the disk. The following values are valid: alphanumeric text strings with a maximum length of 32 characters.

Disk Used (MB) The amount of used space on a disk, expressed in megabytes. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Disk Used Percent The amount of used space on a disk, expressed as a percentage. Note: the value -1 indicates Not Available and the value -2 indicates Not Collected.

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 file system type, such as hfs, nfs, tmpfs, and ufs.

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. If the value for Inodes Free is less than 100, this is a critical condition. Notify your system administrator immediately. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Inodes Free Percent The number of inodes currently available on your file system , expressed as a percentage. Note: the value -1 indicates Not Available and the value -2 indicates Not Collected.

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

Inodes Used Percent The percentage of inodes currently allocated to files, calculated by dividing the Inodes Used value by the Total Inodes value. The following values are valid: integers between 0 and 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. The following values are valid: up to 256 letters or numbers representing a directory path.

Size (MB) The total size of a file system , expressed in megabytes. For example, 1000 represents one gigabyte. The following values are valid: integers of up to 999999999. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

Time Stamp 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 . For example, a value of 163817 indicates that the number of inodes allocated is 163,817. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

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.

Disk attributes (superseded)

The Disk attributes refer to disk characteristics such as inode size, inodes used, mount point, and space available. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Disk Mount Point The path name of the directory to which a file system is mounted. This is the virtual name for the directory. The following values are valid: up to 32 letters or numbers representing a directory path. (Superseded.)

Disk Name The name of the physical disk partition where the file system is mounted. This is the physical location of the disk. The following values are valid: alphanumeric text strings with a maximum length of 32 characters. (Superseded.)

File System Type The file system type, such as *hsfs*, *nfs*, *tmpfs*, and *ufs*. The following values are valid: up to eight letters or numbers. (Superseded.)

Inodes Available Percent The percentage of inodes currently available. Note: the value -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

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. If the value for Inodes Free is less than 100, this is a critical condition. Notify your system administrator immediately. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

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

Inodes Used Percent The percentage of inodes currently allocated to files, calculated by dividing the Inodes Used value by the Total Inodes value. The following values are valid: integers between 0 and 100, such as 85 for 85%. Note: the value -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

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

Size (MB) The total size of a file system , expressed in megabytes. For example, 1000 represents one gigabyte. The following values are valid: integers of up to 99999999. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Space Available (MB) The amount of unused space currently available to non-superusers on a file system , expressed in megabytes. For example, 40000 represents 40 megabytes. The following values are valid: integers of up to 99999999. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

This disk space does not include any space which is reserved for superuser. 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 system s, relative to the disk size, you might need to evaluate reconfiguring the file system to distribute the files more evenly across disks.

Space Available Percent The percentage of space available. The following values are valid: integers between 0 and 100, such as 10 for 10%. Note: the value -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

Space Used (MB) The amount of disk space currently in use on a file system , expressed in megabytes. For example, 5000 represents 5 gigabytes. The following values are valid: integers of up to 99999999. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

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. The following values are

valid: integers between 0 and 100, such as 80 for 80%. Note: the value -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Total Inodes 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, the value -2 indicates Not Collected, and the value 2147483647 indicates *Value_Exceeds_Maximum*. (Superseded.)

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.

Disk IO attributes

The Disk IO attributes refer to disk input/output characteristics, including transfer rates, block read rates, and block write rates.

Note: These attributes are only available for systems with a 2.4 (or higher) kernel.

Blocks Reads Per Second Indicates the amount of data read from the drive expressed in a number of blocks per second. A block is of indeterminate size. The following values are valid: integers. For example, the following value is valid: *Value_Exceeds_Maximum=9223372036854775807*.

Blocks Written Per Second Indicates the amount of data written to the drive expressed in a number of blocks per second. A block is of indeterminate size. The following values are valid: integers. For example, the following value is valid: *Value_Exceeds_Maximum=9223372036854775807*.

Blocks Read The total number of blocks read. The following values are valid: integers. For example, the following value is valid: *Value_Exceeds_Maximum=9223372036854775807*.

Blocks Written The total number of blocks written. The following values are valid: integers. For example, the following value is valid: *Value_Exceeds_Maximum=9223372036854775807*.

Device Major Number Major number of the device. The following values are valid: integers. For example, the following value is valid: *Value_Exceeds_Maximum=2147483647*.

Device Minor Number Distinctive minor number for device. The following values are valid: integers. For example, the following value is valid:
Value_Exceeds_Maximum=2147483647.

Device Name Name of the device as is displayed under the /dev directory.

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

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

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.

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

Transfers Per Second Indicates the number of transfers per second that were issued to the device. A transfer is an I/O request to the device. Multiple logical requests can be combined into a single I/O request to the device. A transfer is of indeterminate size. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Disk IO attributes (superseded)

The Disk IO attributes refer to disk input/output characteristics, including transfer rates, block read rates, and block write rates. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Note: These attributes are only available for systems with a 2.4 (or higher) kernel.

Block Reads Per Second Indicates the amount of data read from the drive expressed in a number of blocks per second. A block is of indeterminate size. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Blocks Written Per Second Indicates the amount of data written to the drive expressed in a number of blocks per second. A block is of indeterminate size. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Blocks Read The total number of blocks read. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Blocks Written The total number of blocks written. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Device Major Number Major number of the device. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Device Minor Number Distinctive minor number for device. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Device Name Name of the device as is displayed under the /dev directory. (Superseded.)

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

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

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Transfers Per Second Indicates the number of transfers per second that were issued to the device. A transfer is an I/O request to the device. Multiple logical requests can be combined into a single I/O request to the device. A transfer is of indeterminate size. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Disk Usage Trends attributes

The Disk Usage Trends attributes refer to disk usage characteristics, such as high water / low water usage rates and days until the disk is full.

Disk Name The name of the physical disk partition where the file system is mounted. This is the physical location of the disk. The following values are valid: alphanumeric text strings with a maximum length of 32 characters.

Days Until Full Disk Current Rate The number of days until the disk is full based on the current rate of disk usage. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum.

Days Until Full Disk Moving Avg The number of days until the disk is full based on the moving average rate of disk usage. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum.

Days Until Full Disk Low Water Mark The number of days until the disk is full based on the disk usage rate that represents the low water mark. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum.

Days Until Full Disk Peak Rate Days until full disk based on the Peak Rate. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum.

Disk Usage Moving Avg (Bytes/Hr) The bytes/hour of disk usage averaged over all previous samples. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`.

Disk Usage Rate (Bytes/Hr) The bytes/hour of disk usage over the last hour. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`.

High Water Mark Disk Usage Rate (Bytes/Hr) The bytes/hour rate that represents the highwater mark of disk usage. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`.

High Water Mark Time Stamp The date and time that the disk usage reaches a highwater mark. The timestamp for SCAN and STR functions is displayed in the CYYMMDDHHMMSSmmm format (as in 1020315064501000 for 03/15/02 06:45:01) where:

Space Free (MB) The amount of unused space currently available to non-superusers on a file system , expressed in megabytes. For example, 40,000 represents 40 megabytes. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`.

This disk space does not include any space which is reserved for superuser. 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 system s, relative to the disk size, you might need to evaluate reconfiguring the file system to distribute the files more evenly across disks.

Space Used (MB) The amount of disk space currently in use on a file system , expressed in megabytes. Valid entries For example, 5000 represents 5 gigabytes. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`.

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

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

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.

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

Disk Usage Trends attributes (superseded)

The Disk Usage Trends attributes refer to disk usage characteristics, such as high water / low water usage rates and days until the disk is full. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Disk Name The name of the physical disk partition where the file system is mounted. This is the physical location of the disk. The following values are valid: alphanumeric text strings with a maximum length of 32 characters. (Superseded.)

Days Until Full Disk Current Rate The number of days until the disk is full based on the current rate of disk usage. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Days Until Full Disk Moving Avg The number of days until the disk is full based on the moving average rate of disk usage. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Days Until Full Disk Low Water Mark The number of days until the disk is full based on the disk usage rate that represents the low water mark. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Days Until Full Disk Peak Rate Days until full disk based on the Peak Rate. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Disk Usage Moving Avg (Bytes/Hr) The bytes/hour of disk usage averaged over all previous samples. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Disk Usage Rate (Bytes/Hr) The bytes/hour of disk usage over the last hour. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

High Water Mark Disk Usage Rate (Bytes/Hr) The bytes/hour rate that represents the highwater mark of disk usage. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Minimum=-2147483648. (Superseded.)

High Water Mark Time Stamp The date and time that the disk usage reaches a highwater mark. (Superseded.)

Space Available (MB) The amount of unused space currently available to non-superusers on a file system , expressed in megabytes. For example, 40,000 represents 40 megabytes. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

This disk space does not include any space which is reserved for superuser. 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 system s, relative to the disk size, you might need to evaluate reconfiguring the file system to distribute the files more evenly across disks.

Space Used (MB) The amount of disk space currently in use on a file system , expressed in megabytes. Valid entries For example, 5000 represents 5 gigabytes.

The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

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

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

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

File Comparison Group attributes

File Comparison Group Attributes refer to File Comparison Group characteristics. This attribute group is not available for historical data collection.

File Compare Option The File compare option is used to specify which type of comparison is used. The following values are valid: Plain (1), Ignore_Whitespace (2), Ignore_Case (3), Ignore_Case_Whitespace (4), and Binary (5). The default is Plain.

File Compare Result The result of the file comparison between File_Name_1 and File_Name_2. The following values are valid: Same (0) and Different (1). Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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

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

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

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

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.

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

File Information attributes

The File Information attributes refer to file information characteristics. This attribute group is not available for historical data collection.

Access The access rights of the file expressed as 4-digit octal number.

Attribute Last Change Time The date and time of the last file attributes change.

Checksum Checksum or hash string based on hashing algorithm. The default algorithm is CRC32.

Checksum Algorithm Only used in situations in conjunction with the Checksum attribute to specify the algorithm to be used to calculate the hash string. Note: -1 indicates Not_Applicable. Other possible values are CRC32 (0), MD5 (1), and SHA1 (2), Not_Available (-1). The default is CRC32.

File The name of file or directory. If the file is a symbolic link, the link name is shown in Link_Name attribute.

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 (0), Yes (1), and Not Available (-1).

File Mode Mode is the string representation of the access rights of the file. This is related to the Access attribute. The access attribute is the octal representation of the access rights of the file. The mode of a file would be rwxr-xr-x if the access was 755.

Group The logical group to which the file belongs.

Last Accessed Time The date and time of the last file access.

Last Changed Time The date and time of the last change to a file.

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.

Links The number of links to a file.

Owner The name of the file owner.

Path The full path containing a particular file or directory.

Size (MB) The size, in MB, of the file. This attribute displays as a floating point with a scale of 3. For example 55.255. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Size (MB) (Superseded) The size, in MB, of the file. This attribute displays as a floating point with a scale of 3. For example 55.255.

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

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

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.

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

Type The type of file, including Dir (= directory), DirLink (=directory link), File (= file), FileLink (=file link), Sock (= socket), Link (= link), Spec (= special file), and Unknown (=unknown).

File Pattern Group attributes

The File Pattern Group attributes refer to file pattern group characteristics. This attribute group is not available for historical data collection.

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: -1 indicates Not_Available and -2 indicates Not_Collected.

Match Option Options that affect how the search is performed. The following values are valid: Normal (1), Ignore_Case (2), Inverse_Search (3), and Match_Whole_Words (4).

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

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

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

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.

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

I/O Ext attributes

The I/O Ext attributes refer to a wide variety of disk input/output characteristics, including read request rates, write request rates, and service time measures.

Note: These attributes are only available for systems with a 2.4 (or higher) kernel.

Average Request Queue Length The average queue length of the requests that were issued to the device. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Average Request Size (Sectors) The average size (in sectors) of the requests that were issued to the device. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Average Service time (ms) The average service time (in milliseconds) for I/O requests that were issued to the device. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Average Wait Time (ms) The average time (in milliseconds) for I/O requests issued to the device to be served. The following values are valid: integers. For

example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Bytes Transferred Per Second The number of bytes transferred per second. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Device Name Name of the device as is displayed under the /dev directory. The following values are valid: alphanumeric text strings with a maximum length of 64 characters.

Disk Read Percent The percentage of time spent in read operations.

Disk Write Percent The percentage of time spent in write operations.

Percent CPU Time Used Percentage of CPU time during which I/O requests were issued to the device. Saturation occurs at 100%.

Read Bytes Per Second The number of bytes read from the device per second. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Read Requests Per Second The number of read requests that were issued, per second, to the device. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Read Requests Merged Per Second The number of read requests merged, per second, that were issued to the device. The following values are valid: integers. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Read Sectors Per Second The number of sectors read, per second, from the device. The following values are valid: integers. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

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

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

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.

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

Write Bytes Per Second The number of bytes written to the device per second. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Write Requests Per Second The number of write requests that were issued, per second, to the device. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Write Requests Merged Per Second The number of write requests merged that were issued, per second, to the device. The following values are valid: integers. For example, the following value is valid:

Value_Exceeds_Maximum=9223372036854775807.

Write Sectors Per Second The number of sectors written to the device, per second. The following values are valid: integers. For example, the following value is valid:

Value_Exceeds_Maximum=9223372036854775807.

I/O Ext attributes (superseded)

The I/O Ext attributes refer to a wide variety of disk input/output characteristics, including read request rates, write request rates, and service time measures. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Note: These attributes are only available for systems with a 2.4 (or higher) kernel.

Average Request Queue Length The average queue length of the requests that were issued to the device. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Average Request Size (Sectors) The average size (in sectors) of the requests that were issued to the device. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Average Service time (ms) The average service time (in milliseconds) for I/O requests that were issued to the device. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Average Wait Time (ms) The average time (in milliseconds) for I/O requests issued to the device to be served. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Bytes Transferred Per Second The number of bytes transferred per second. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Device Name Name of the device as is displayed under the /dev directory. The following values are valid: alphanumeric text strings with a maximum length of 64 characters. (Superseded.)

Disk Read Percent The percentage of time spent in read operations. (Superseded.)

Disk Write Percent The percentage of time spent in write operations. (Superseded.)

Percent CPU Time Used Percentage of CPU time during which I/O requests were issued to the device. Saturation occurs at 100%. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Read Bytes Per Second The number of bytes read from the device per second. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Read Requests Per Second The number of read requests that were issued, per second, to the device. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Read Requests Merged Per Second The number of read requests merged, per second, that were issued to the device. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Read Sectors Per Second The number of sectors read, per second, from the device. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

System Name The managed system name. The form should be `hostname:agent_code`. (Superseded.)

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

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Write Bytes Per Second The number of bytes written to the device per second. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Write Requests Per Second The number of write requests that were issued, per second, to the device. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Write Requests Merged Per Second The number of write requests merged that were issued, per second, to the device. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Write Sectors Per Second The number of sectors written to the device, per second. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

IP Address attributes

The IP Address attributes refer to network characteristics, including IP address and network interface name.

DNS Name The Domain Name Server (DNS) entry associated with the IP network address. The following values are valid: alphanumeric text strings with a maximum length of 384 characters. Note that the value `No_DNS_Entry` indicates `NO_DNS_ENTRY`.

IP Address An IP address associated with the network interface. The following values are valid: alphanumeric text strings with a maximum length of 46 characters.

IP Version An indicator as to whether the IP address is version 4 or version 6, including IPv4=4 and IPv6=6.

Network Interface Name The name of the network interface. The following values are valid: alphanumeric text strings with a maximum length of 32 characters.

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

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

Linux Group attributes

The Linux Group attributes refer to group characteristics.

Duplicate Group Name True if the group name is listed more than once in `/etc/group`. The following values are valid: False (0) and True (1).

Group ID The ID of this group. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`.

Group ID (Superseded) The ID of this group. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`.

Group Name The name of the group.

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

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

Linux Host Availability attributes

The Linux Host Availability attributes refer to Linux host availability characteristics. The attributes in this group can only be used in a situation. Historical information is available for the Host Availability 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 - "KLZ_PINGHOSTLIST" in the `lz.ini` file in the IBM Tivoli Monitoring config directory. For example:

```
KLZ_PINGHOSTLIST='/opt/ibm/itm/config/klzpinghosts'
```

```
sample content of klzpinghosts:  
#  
# hosts pinged for availability from this agent  
#  
server1.domain.com  
server2  
server4
```

Host Status Result of the "ping" operation. The following values are valid: Successful (1), Unsuccessful (0), and Error (-1).

Server Response Time Ping operation response time in milliseconds. Note: -1000 indicates Not Available.

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

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

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.

Target Host The host name or IP Address of the target of the ping operation. The following values are valid: alphanumeric text strings with a maximum length of 128 characters.

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

Linux TCP Statistics attributes

The Linux 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 KLZ_TCPSTAT_SAMPLE_SECS environment variable. The default value is 30 seconds, and the minimum value is 5 seconds.

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

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

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.

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.

LPAR attributes

The LPAR attributes refer to LPAR characteristics such as entitlement, capacity weight, and shared pool idle time.

Capacity Weight The priority of the LPAR when running in uncapped mode. Used to prioritize partitions competing for CPU resources. If zero is specified, this essentially sets the partition to capped. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Capped Specifies whether the LPAR can exceed or not its entitled capacity. Note: the value 0 indicates false, the value 1 indicates true, the value -1 indicates Not Available, and the value -2 indicates Not Collected.

Entitlement The entitled capacity of the LPAR in terms of units of physical processors. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Entitlement Used (Percent) The percentage of the entitled CPU Units that are in use. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

LPAR ID The identifier of the LPAR. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

LPAR Name The name of the LPAR.

Max Capacity Used (Percent) The percentage of the max capacity units that are in use. For capped LPARs, this value is equal to Entitlement Used (Percent). Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Maximum Physical CPUs The maximum number of physical CPUs in the system. Note: the value -1 indicates Not Available, -2 indicates Not Collected, -2147483648 indicates Value Exceeds Minimum, and 2147483647 indicates Value Exceeds Maximum.

Number of Physical CPUs The number of active licensed physical CPUs. Note: the value -1 indicates Not Available, -2 indicates Not Collected, -2147483648 indicates Value Exceeds Minimum, and 2147483647 indicates Value Exceeds Maximum.

Number of Virtual CPUs The number of virtual processors of the LPAR. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Physical CPU Units Used The average number of physical CPU units used in the last sampling interval. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and 9223372036854775807 indicates Value Exceeds Maximum.

Shared Specifies whether the LPAR is configured for sharing or not-sharing its free CPU resources with other LPARs. Note: the value 0 indicates false, the value 1 indicates true, the value -1 indicates Not Available, and the value -2 indicates Not Collected.

Shared Pool Idle Time The average number of CPU units left unused in the shared processor pool over the last sampling interval. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and 9223372036854775807 indicates Value Exceeds Maximum.

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

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

Total Steal Time (Percent) The total Steal time reported by all the virtual CPUs defined in the LPAR, as percent over the last sampling interval. Note: the value -1 indicates Not Available, and the value -2 indicates Not Collected. Note that the attribute value is averaged in synch with the situation or historical collection interval.

Machine Information attributes

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

Note: This Monitoring Agent for Linux gathers the following attributes in this group by using the command `/usr/sbin/dmidecode`:

- BIOS Version
- BIOS Release
- Hardware Brand
- Hardware Model
- Machine Serial Number

The Monitoring Agent for Linux must be running as root in order to execute this command. If not, "Unknown" is returned for the dmidecode metrics. Further, this program is not available for zLinux or pSeries systems. Metrics other than Hardware Brand, Hardware Model and Machine Serial Number will report as "Unknown." Further information on dmidecode is available at the following website:

<http://www.nongnu.org/dmidecode>

BIOS Release The BIOS vendor release date. Note: the value unknown = UNKNOWN.

BIOS Version The BIOS vendor version. Note: the value unknown = UNKNOWN.

Hardware Brand The brand of hardware on which the agent is running. Note: the value unknown = UNKNOWN.

Hardware Model The specific hardware model underlying the monitored operating system. Note: the value unknown = UNKNOWN.

Machine Serial Number The serial number of the computer. Note: the value unknown = UNKNOWN.

Number of Processors Configured The number of processors configured for this computer. This number excludes secondary processor contexts, but might include virtual processors in some virtual environments. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Number of Processors Online The number of processors online the computer. This number excludes secondary processor contexts, but might include virtual processors in some virtual environments. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Server Host Name The host name for the computer. Note: the value unknown = UNKNOWN.

System Board UUID The Universally Unique Identifier burned in to the system board.

System Name The managed system name. The form should be *hostname:agent_code*. Note: the value unknown = UNKNOWN.

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

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

Virtual Machine Identifier The identifier of the operating system instance in a virtualized environment (zLinux and Linux PPC only).

Network attributes

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

Bytes Received Per Second The number of bytes received per second by the interface. The following values are valid: integers in the range 0 to 9223372036854775807. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`. The value -1 indicates Not_Available.

Bytes Transmitted Per Second The number of bytes transmitted per second by the interface. The following values are valid: integers in the range 0 to 9223372036854775807. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`. The value -1 indicates Not_Available.

Carrier Losses The number of carrier losses that occurred in the interface. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`.

Collisions (Percent) Of the total number of packets transmitted in this sample period, the percentage involved in a collision. The following values are valid: integers.

Collisions Per Minute The number of times a packet collided with another packet per minute. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807. The value -1 indicates Not_Available.

Device Type The device type. The following values are valid: NETROM, ETHER, EETHER, AX25, PRONET, CHAOS, IEEE802_TR, ARCNET, APPLETLK, DLCI, ATM, METRICOM, IEEE1394, SLIP, CSLIP, SLIP6, CSLIP6, RSRVD, ADAPT, ROSE, X25, HWX25, PPP, HDLC, LAPB, DDCMP, RAWHDLC, TUNNEL, TUNNEL6, FRAD, SKIP, LOOPBACK, LOCALTLK, FDDI, BIF, SIT, IPDDP, IPGRE, PIMREG, HIPPI, ASH, ECONET, IRDA, FCPP, FCAL, FCPL, FCFABRIC, IEEE802, IEEE80211, UNKNOWN.

Errors (Percent) Of the total number of packets received and transmitted, the percentage that were in error during this sample period. The following values are valid: integers.

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

Input Error (Percent) The number of input packet errors as a percentage of the total number of packets received in this sample.

Input Errors The number of packets with errors received on the interface. Valid entry is an integer in the range zero to 9223372036854775807. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807. The value -1 indicates Not_Available.

Input Errors Per Minute The number of packets with errors received per minute by the interface. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Input FIFO Buffer Overruns The number of input FIFO buffer overruns that occurred during the sampling period. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Input Packets Dropped The number of input packets dropped by the device driver. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Example: www.company.com indicates that the DNS will resolve the name www.company.com to mean the IP address for the interface.

IPv4 Address The Internet Protocol (IP) address of the network interface. A gateway system might have more than one interface, each with a separate address. The following values are valid: 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.

Interface Status This attribute indicates if a network interface is currently available. Valid entries include 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), and UNKNOWN (indicates the interface is in unknown). Note that these values are case-sensitive. Example:UP means an interface is in service.

MAC Address The MAC address of the Network Interface Card. NOT_AVAILABLE is a valid value. It is typically 6 bytes, but can be up to 14. The value is formatted with a colon between each byte.

Maximum Transmission Unit The maximum packet size (in bytes) for the specified network interface. This is a fixed value. The following values are valid: integers in the range 0 to 99999999. 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.

Network Interface Name Identifies the network interface adapter. The following values are valid: 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 0 to 7.

Output Errors The number of packet transmission errors by the network interface. Valid entry is an integer in the range zero to 9223372036854775807. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Output Error (Percent) The total number of output errors as a percentage of the total number of packets transmitted in this sample.

Output Errors Per Minute The number of packet transmission errors per minute during the monitoring interval. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807. The value -1 indicates Not_Available.

Output FIFO Buffer Overruns The number of output FIFO buffer overruns that occurred during the sampling period. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Output Packets Dropped The number of output packets dropped by the device driver. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Packet Framing Errors The number of packet framing errors that occurred in the interface. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Packets Received The number of packets received by the interface. Valid entry is an integer in the range zero to 9223372036854775807. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Packets Received Per Second The number of packets received per second by the interface during the sampling period. Valid entry is an integer in the range zero to

9223372036854775807. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807. The value -1 indicates
Not_Available.

Packets Transmitted The number of packets transmitted by the interface. Valid entry is an integer in the range zero to 9223372036854775807. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Packets Transmitted Per Second The number of packets transmitted per second by the interface during the sampling period. Valid entry is an integer in the range zero to 9223372036854775807. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807. The value -1 indicates Not_Available.

Received Count (KB) The number of kilobytes received since the network interface was configured. Valid entry is an integer in the range zero to 9223372036854775807. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

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

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

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

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.

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

Total Collisions The number of times during the sampling period 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. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

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.

Transmitted Count (KB) The number of kilobytes transmitted by an interface since boot time. Valid entry is an integer in the range zero to 9223372036854775807. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

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.

Network attributes (superseded)

The Network attributes refer to network characteristics such as received count, sent count, network interface name, and interface status. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Bytes Received Per Second The number of bytes received per second by the interface. The following values are valid: integers in the range 0 to 2147483647. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Bytes Transmitted Per Second The number of bytes transmitted per second by the interface. The following values are valid: integers in the range 0 to 2147483647. For example, the following value is valid: Value_Exceeds_Minimum=-2147483648 and the value Value_Exceeds_Maximum=2147483647. (Superseded.)

Carrier Losses The number of carrier losses that occurred in the interface. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Collisions (Percent) Of the total number of packets transmitted in this sample period, the percentage involved in a collision. The following values are valid: integers. (Superseded.)

Collisions Per Minute The number of times a packet collided with another packet per minute. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Device Type The device type. The following values are valid: NETROM, ETHER, EETHER, AX25, PRONET, CHAOS, IEEE802_TR, ARCNET, APPLETLK, DLCI, ATM, METRICOM, IEEE1394, SLIP, CSLIP, SLIP6, CSLIP6, RSRVD, ADAPT, ROSE, X25, HWX25, PPP, HDLC, LAPB, DDCMP, RAWHDLC, TUNNEL, TUNNEL6, FRAD, SKIP, LOOPBACK, LOCALTLK, FDDI, BIF, SIT, IPDDP, IPGRE, PIMREG, HIPPI, ASH, ECONET, IRDA, FCPP, FCAL, FCPL, FCFABRIC, IEEE802, IEEE80211, UNKNOWN. (Superseded.)

Errors (Percent) Of the total number of packets received and transmitted, the percentage that were in error during this sample period. The following values are valid: integers. (Superseded.)

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

Input Error (Percent) The number of input packet errors as a percentage of the total number of packets received in this sample. (Superseded.)

Input Errors The number of packets with errors received on the interface. The following values are valid: integers in the range 0 to 100. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Input Errors Per Minute The number of packets with errors received per minute by the interface. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Input FIFO Buffer Overruns The number of input FIFO buffer overruns that occurred during the sampling period. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Input Packets Dropped The number of input packets dropped by the device driver. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Example: www.company.com indicates that the DNS will resolve the name www.company.com to mean the IP address for the interface.

IPv4 Address The Internet Protocol (IP) address of the network interface. A gateway system might have more than one interface, each with a separate address. The following values are valid: Internet protocol addresses in the form a.b.c.d. where a, b, c, and d are integers in the range 0 to 255. (Superseded.)

Example: 197.128.55.55 indicates the network interface uses the IP address 197.128.55.55.

Interface Status This attribute indicates if a network interface is currently available. (Superseded.) Valid entries include 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), and UNKNOWN (indicates the interface is in unknown). Note that these values are case-sensitive. Example:UP means an interface is in service.

MAC Address The MAC address of the Network Interface Card. NOT_AVAILABLE is a valid value. It is typically 6 bytes, but can be up to 14. The value is formatted with a colon between each byte. (Superseded.)

Maximum Transmission Unit The maximum packet size (in bytes) for the specified network interface. This is a fixed value. The following values are valid: integers in the range 0 to 99999999. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. 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. (Superseded.)

Network Interface Name Identifies the network interface adapter. (Superseded.) The following values are valid: 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 0 to 7.

Output Errors The number of packet transmission errors by the network interface. The following values are valid: integers in the range 0 to 100. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Output Error (Percent) The total number of output errors as a percentage of the total number of packets transmitted in this sample. (Superseded.)

Output Errors Per Minute The number of packet transmission errors per minute during the monitoring interval. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Output FIFO Buffer Overruns The number of output FIFO buffer overruns that occurred during the sampling period. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Output Packets Dropped The number of output packets dropped by the device driver. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Packet Framing Errors The number of packet framing errors that occurred in the interface. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Packets Received The number of packets received by the interface. The following values are valid: integers in the range 0 to 99999999. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Packets Received Per Second The number of packets received per second by the interface during the sampling period. The following values are valid: integers in the range 0 to 2147483647. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Packets Transmitted The number of packets transmitted by the interface. The following values are valid: integers in the range 0 to 99999999. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Packets Transmitted Per Second The number of packets transmitted per second by the interface during the sampling period. The following values are valid: integers in the range 0 to 2147483647. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Received Count (KB) The number of kilobytes received since the network interface was configured. The following values are valid: integers in the range 0 to 2147483647. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

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

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

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

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Total Collisions The number of times during the sampling period 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. The following values are valid: integers in the range 0 to 100. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

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.

Transmitted Count (KB) The number of kilobytes transmitted by an interface since boot time. The following values are valid: integers in the range 0 to 2147483647. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

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.

NFS Statistics attributes

Use NFS Statistics to monitor characteristics of Network File System (NFS) such as the number of calls, lookups, and operations. This agent currently reports only on NFS version 2 and 3 statistics.

Access Calls The number of access calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Access Calls Percent Of the total number of calls made to the NFS server, the percentage that were access calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Commit Calls The number of file commit calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Commit Calls Percent Of the total number of calls made to the NFS server, the percentage that were file commit calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

File Creates The number of file create calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

File Creates Percent Of the total number of calls made to the NFS server, the percentage that contained file creation operations. The following values are valid: integers in the range of 0 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

File System Info Calls The number of file system information calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

File System Info Calls Percent Of the total number of calls made to the NFS server, the percentage that were calls to obtain information about the file system. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

File System Statistics Calls The number of calls made to the NFS server which requested statistics of the file system. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

File System Statistics Calls Percent Of the total number of calls made to the NFS server, the percentage that involved a request for file system statistics. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Get Attribute Calls The number of calls made to the NFS server which contained a get attribute (getattr) operation. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Get Attribute Calls Percent Of the total number of calls made to the NFS server, the percentage that contained get attribute (getattr) operations. The following values are valid: integers in the range of 0 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Link Calls The total number of link calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Link Calls Percent Of the total number of calls made to the NFS server, the percentage that were link calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Location The location of the origin of the call in the Network File System. The following values are valid: integers. A value of 0 indicates unknown, the value of 1 represents the server, and a value of 2 represents the client. Note: the value -1 indicates Not Available and the value -2 indicates Not Collected.

Lookups The number of lookups made on the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Lookups Percent Of the total number of calls made to the NFS server, the percentage that were lookups. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Make Directory Calls The number of make directory calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Make Directory Calls Percent Of the total number of calls made to the NFS server, the percentage that were make directory calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Make Node Calls The number of make node (mknod) calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Make Node Calls Percent Of the total number of calls made to the NFS server, the percentage that were make node (mknod) calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

NFS Calls The total NFS server or client calls. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

NFS Version The software version associated with the NFS server. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Null Calls The number of calls made to the NFS server from NFS clients which contained no data. The following values are valid: integers in the range of 0 to 100. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Null Calls Percent Of the total number of calls made to the NFS server, the percentage that contained no data. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Path Conf Calls The number of calls made to the NFS server that involved path configuration (pathconf) calls to obtain configuration values for files. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Path Conf Call Percent Of the total number of calls made to the NFS server, the percentage that involved use of the pathconf command to obtain configuration values for files. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Read Calls The number of read calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Read Calls Percent Of the total number of calls made to the NFS server, the percentage that were read calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Read Dir Plus Calls The number of read directory plus (readdirplus) calls made to the NFS server to return the name, the file ID, attributes, and file handle. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Read Dir Plus Calls Percent Of the total number of calls made to the NFS server, the percentage that were read directory plus (readdirplus) calls. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Read Directory Calls The number of read directory calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Read Directory Calls Percent Of the total number of calls made to the NFS server, the percentage that were read directory calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Read Link Calls The number of read link calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Read Link Calls Percent Of the total number of calls made to the NFS server, the percentage that were read link calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Remove Directory Calls The number of remove directory calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Remove Directory Calls Percent Of the total number of calls made to the NFS server, the percentage that were remove directory calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Remove File Calls The number of file removal calls made to the NFS server. The following values are valid: integers. Note: -1 indicates Not_Available, -2 indicates Not_Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum..

Remove File Calls Percent Of the total number of calls made to the NFS server, the percentage that were file removal calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Rename File Calls The number of file rename calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Rename File Calls Percent Of the total number of calls made to the NFS server, the percentage that were file rename calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Root Calls The number of calls made to the NFS server which contained root calls. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Root Calls Percent Of the total number of calls made to the NFS server, the percentage that were root calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Set Attribute Calls The number of calls made to the NFS server which contained a set attribute (setattr) operation. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Set Attribute Calls Percent Of the total number of calls made to the NFS server, the percentage that contained a set attribute (setattr) operation. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Symbolic Link Calls The total number of symbolic link calls. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Symbolic Link Calls Percentage Of the total number of calls made to the NFS server, the percentage that were symbol link calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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

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

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.

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

Write Cache Calls The number of write cache calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Write Cache Calls Percent Of the total number of calls made to the NFS server, the percentage that were write cache calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Writes The number of write calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Writes Percent Of the total number of calls made to the NFS server, the percentage that were write calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

NFS Statistics attributes (superseded)

Use NFS Statistics to monitor characteristics of Network File System (NFS) such as the number of calls, lookups, and operations. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Access Calls The number of access calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Access Calls Percent Of the total number of calls made to the NFS server, the percentage that were access calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Commit Calls The number of file commit calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Commit Calls Percent Of the total number of calls made to the NFS server, the percentage that were file commit calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

File Creates The number of file create calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

File Creates Percent Of the total number of calls made to the NFS server, the percentage that contained file creation operations. The following values are valid: integers in the range of 0 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

File System Info Calls The number of file system information calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

File System Info Calls Percent Of the total number of calls made to the NFS server, the percentage that were calls to obtain information about the file system. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

File System Statistics Calls The number of calls made to the NFS server which requested statistics of the file system. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

File System Statistics Calls Percent Of the total number of calls made to the NFS server, the percentage that involved a request for file system statistics. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Get Attribute Calls The number of calls made to the NFS server which contained a get attribute (getattr) operation. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Get Attribute Calls Percent Of the total number of calls made to the NFS server, the percentage that contained get attribute (getattr) operations. The following values are valid: integers in the range of 0 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Link Calls The total number of link calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Link Calls Percent Of the total number of calls made to the NFS server, the percentage that were link calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Location The location of the origin of the call in the Network File System. The following values are valid: integers. A value of 0 indicates unknown, the value of 1 represents the server, and a value of 2 represents the client. Note: the value -1 indicates Not Available and the value -2 indicates Not Collected. (Superseded.)

Lookups The number of lookups made on the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Lookups Percent Of the total number of calls made to the NFS server, the percentage that were lookups. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Make Directory Calls The number of make directory calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Make Directory Calls Percent Of the total number of calls made to the NFS server, the percentage that were make directory calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Make Node Calls The number of make node (mknod) calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Make Node Calls Percent Of the total number of calls made to the NFS server, the percentage that were make node (mknod) calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

NFS Calls The total NFS server or client calls. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

NFS Version The software version associated with the NFS server. The following values are valid: integers. A value of 2 represents version 2, 3 represents version 3, 4 represents version 4. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Null Calls The number of calls made to the NFS server from NFS clients which contained no data. The following values are valid: integers in the range of 0 to 100. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Null Calls Percent Of the total number of calls made to the NFS server, the percentage that contained no data. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Path Conf Calls The number of calls made to the NFS server that involved path configuration (pathconf) calls to obtain configuration values for files. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Path Conf Call Percent Of the total number of calls made to the NFS server, the percentage that involved use of the pathconf command to obtain configuration values for files. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Read Calls The number of read calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Read Calls Percent Of the total number of calls made to the NFS server, the percentage that were read calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Read Dir Plus Calls The number of read directory plus (readdirplus) calls made to the NFS server to return the name, the file ID, attributes, and file handle. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Read Dir Plus Calls Percent Of the total number of calls made to the NFS server, the percentage that were read directory plus (readdirplus) calls. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Read Directory Calls The number of read directory calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Read Directory Calls Percent Of the total number of calls made to the NFS server, the percentage that were read directory calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Read Link Calls The number of read link calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Read Link Calls Percent Of the total number of calls made to the NFS server, the percentage that were read link calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Remove Directory Calls The number of remove directory calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Remove Directory Calls Percent Of the total number of calls made to the NFS server, the percentage that were remove directory calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Remove File Calls The number of file removal calls made to the NFS server. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Remove File Calls Percent Of the total number of calls made to the NFS server, the percentage that were file removal calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Rename File Calls The number of file rename calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Rename File Calls Percent Of the total number of calls made to the NFS server, the percentage that were file rename calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Root Calls The number of calls made to the NFS server which contained root calls. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Root Calls Percent Of the total number of calls made to the NFS server, the percentage that were root calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Set Attribute Calls The number of calls made to the NFS server which contained a set attribute (setattr) operation. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Set Attribute Calls Percent Of the total number of calls made to the NFS server, the percentage that contained a set attribute (setattr) operation. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Symbolic Link Calls The total number of symbolic link calls. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Symbolic Link Calls Percentage Of the total number of calls made to the NFS server, the percentage that were symbol link calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

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

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

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Write Cache Calls The number of write cache calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Write Cache Calls Percent Of the total number of calls made to the NFS server, the percentage that were write cache calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Writes The number of write calls made to the NFS server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Writes Percent Of the total number of calls made to the NFS server, the percentage that were write calls. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

OS Configuration attributes

The OS Configuration attributes refer to configuration characteristics such as OS Name and OS Version.

GCC Version The version of the GNU Compiler with which the kernel was compiled.

OS Name The operating system name.

OS Vendor Information The operating system information.

OS Version The operating system version.

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

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

Vendor ID The Processor Vendor ID.

Process attributes

The Process attributes refer to process characteristics such as data set size, kernel scheduling priority, the number of pages of memory, and the number of page faults. The sampling interval for the Process Instant Busy CPU (Percent) attribute is configured using the `KLZ_PROCESS_SAMPLE_SECS` environment variable. The default value is 30 seconds. The value 0 specifies that the Instant Busy CPU calculation is disabled.

Command Line The process command line string. The following values are valid: text strings with a maximum length of 768 characters.

CPU Seconds The total CPU seconds spent on the process. The following is a valid value: `Value_Exceeds_Maximum=9223372036854775807`.

Cumulative Busy CPU (Percent) The summation of user CPU and system CPU for this process's children.

Cumulative Process System CPU (Percent) The percentage of cumulative CPU time spent in kernel mode by this process's children. The following values are valid: integers between 0 and 100.

Cumulative Process User CPU (Percent) The percentage of cumulative CPU time spent in user mode by this process's children. The following values are valid: integers between 0 and 100.

Data Resident Set (Pages) The size of the data set based on the number of pages. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807` and `Not_Collected=-2`.

Data Size (KB) The data size (in kilobytes) of the virtual memory. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807 Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Data Size (MB) The data size (in megabytes) of the virtual memory. This attribute displays as a floating point with a scale of 1. For example 5.2. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Executable Size (KB) The executable size (in kilobytes) of the virtual memory. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Executable Size (MB) The executable size (in megabytes) of the virtual memory. This attribute displays as a floating point with a scale of 1. For example 5.2. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Kernel Priority The kernel scheduling priority. The following values are valid: integers between -100 - 100 (-100 is the highest). Real-time processes can have priorities that are negative.

Library Size (KB) The library size (in kilobytes) of the virtual memory. This measurement represents all pages, including unused. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Library Size (MB) The library size (in megabytes) of the virtual memory. This measurement represents all pages, including unused. This attribute displays as a floating point with a scale of 1. For example 5.2. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Nice Value The standard UNIX nice level (-20 represents the highest level). The following values are valid: integers in the range -20 to 19.

Number of Threads The number of threads started for this process. (Valid only on 2.6 kernel and above.) Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Process Busy CPU (Percent) The summation of User CPU Percent and System CPU Percent for this process.

Process Busy CPU (Time) The total CPU time (user + system) spent on the process. The value is expressed in the DDDd HHh MMm SSs format. The following values are valid: Not_Available=0.

Process Command Name The name of the process command. The following values are valid: alphanumeric text strings with a maximum length of 96 characters.

Process Count The count of processes with the same name. The name is selected by using the Command Line (UNICODE), CMDLINEU, attribute. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Process CPU ID The ID of the process CPU. The following values are valid: integers. Note: -1 indicates Not Available.

Process Dirty Pages Pages that have been modified (dirty) in buffer (main memory), but not yet copied to the cache. The following values are valid: integers. For example, the following values are valid:
Value_Exceeds_Maximum=9223372036854775807 and Not_Collected=-2.

Process Filter A regular expression to be applied to the Command Line attribute. The maximum allowable length is 256 characters. Valid values include
Java_processes_(.*java.*),
IBM_Java_processes_entry_method_only_(.*java.*(com.ibm.*)), and
System_Admin_installed_processes_(/usr.*).

Process Group LeaderID The process group leader ID. The following values are valid: integers. For example, the following values are valid:
Value_Exceeds_Maximum=9223372036854775807, Not Available=-1, and
Not_Collected=-2.

Process ID The identifier of the process. The following values are valid: integers between 0 and 999. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Process Instant Busy CPU (Percent) The percentage of CPU used by this process in the last sampling interval normalized to account for multiple online processors. For example, the following values are valid: Not Available=-1, and
Not_Collected=-2. Note that the attribute value is averaged in synch with the situation or historical collection interval.

Process Parent ID The identifier for the parent process. The following values are valid: integers between 0 and 999. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Process Short Term Avg Busy CPU (Percent) The summation of Proc System CPU Norm and Proc User CPU Norm for this process. CPU percentages are normalized to account for multiple online processors; percentages are normalized to a maximum of 100 percent. This metric is only available through situations and only when the Process name is also specified within the situation predicate. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Process Short Term Avg System CPU (Percent) The short term average of the percentage of CPU time spent in kernel mode by the process. The attribute value, comparable to the output of the **top -s** command, is taken on the interval of 30 seconds. CPU percentages are normalized to account for multiple online processors; percentages are normalized to a maximum of 100 percent. This metric is only available through situations and only when the Process name is also specified within the situation predicate. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Process Short Term Avg User CPU (Percent) The short term average of the percentage of CPU time spent in user mode by the process. The attribute value, comparable to the output of the **top -s** command, is taken on the interval of 30 seconds. CPU percentages are normalized to account for multiple online processors; percentages are normalized to a maximum of 100 percent. This metric

is only available through situations and only when the Process name is also specified within the situation predicate. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Process State The state of the process (Sleeping, Disk, Running, Zombie, Trace, Dead, or N/A). The following values are valid: integers between -1 and 5, where 0 = Sleeping, 1 = Disk, 2 = Running, 3 = Zombie, 4 = Trace, 5 = Dead, and -1 = Not_Available.

Process System CPU (Percent) The percentage of CPU time spent in kernel mode by process. The attribute value, comparable with the output of the **ps -ef** command, is determined by the system time of the process divided by the time since boot. The following values are valid: integers between 0 and 100.

Process System CPU (Time) The system CPU time spent executing the process. The value is expressed in the DDDd HHh MMm SSs format. The following values are valid: Not_Available=0.

Process User CPU (Percent) The percentage of CPU time spent in user mode by process. The attribute value, comparable with the output of the **ps -ef** command, is determined by the user time of the process divided by the time since boot. The following values are valid: integers between 0 and 100.

Process User CPU (Time) The user CPU time spent executing the process. The value is expressed in the DDDd HHh MMm SSs format. The following values are valid: Not_Available=0.

Resident Set Size (Pages) The number of pages the process has in real memory. The following values are valid: integers. For example, the following values are valid: Value_Exceeds_Maximum=9223372036854775807 and Not_Collected=-2.

Session ID The session ID. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Shared Lib Resident Set (Pages) The number of pages of shared library set (mmap) memory. mmap is a system API that lets you map a file or device into memory. The mapped pages might be shared so that other processes can access them. The following values are valid: integers. For example, the following values are valid: Value_Exceeds_Maximum=9223372036854775807 and Not_Collected=-2.

Shared Memory (Pages) The number of pages of shared (mmap) memory. mmap is a system API that lets you map a file or device into memory. The mapped pages might be shared so that other processes can access them. The following values are valid: integers. For example, the following values are valid: Value_Exceeds_Maximum=9223372036854775807 and Not_Collected=-2.

Stack Size (KB) The stack size (in kilobytes) of the virtual memory. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Stack Size (MB) The stack size (in megabytes) of the virtual memory. The following values are valid: integers. This attribute displays as a floating point with

a scale of 1. For example 5.2. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

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

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

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.

Text Resident Set (Pages) The number of pages of text resident (mmap) memory. mmap is a system API that lets you map a file or device into memory. The mapped pages might be shared so that other processes can access them. The following values are valid: integers. For example, the following values are valid: Value_Exceeds_Maximum=9223372036854775807 and Not_Collected=-2.

Time The total amount of CPU time that a process has consumed.

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

Total Major Faults The total number of major page faults (including child processes) since the start of the process. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Total Minor Faults The total number of minor page faults (including child processes) since the start of the process. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Total Size (Pages) The number of pages that the process has in real memory. The following values are valid: integers. For example, the following values are valid: Value_Exceeds_Maximum=9223372036854775807 and Not_Collected=-2.

User to System CPU (Percent) Of the total system CPU usage, the percentage that was user CPU usage. For example, 500% means that user CPU usage is 5 times the system CPU usage. The following values are valid: integers between -10,000 and 10,000.

VM Locked Pages (KB) The size (in kilobytes) of locked pages of the virtual memory. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

VM Locked Pages (MB) The size (in megabytes) of locked pages of the virtual memory. This attribute displays as a floating point with a scale of 1. For example 5.2. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

VM Size (KB) The size (in kilobytes) of the virtual memory. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

VM Size MB Virtual memory size in megabytes. This attribute displays as a floating point with a scale of 1. For example 5.2. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Process attributes (superseded)

The Process attributes refer to process characteristics such as data set size, kernel scheduling priority, the number of pages of memory, and the number of page faults. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Command Line The process command line string. The following values are valid: alphanumeric text strings with a maximum length of 256 characters. (Superseded.)

Command Line (Unicode) The process command line string. The following values are valid: text strings with a maximum length of 512 bytes. This attribute is globalized (Unicode). (Superseded.)

Cumulative Busy CPU (Percent) The summation of user CPU and system CPU for this process and children. (Superseded.)

Cumulative Process System CPU (Percent) The percentage of cumulative CPU time spent in kernel mode by process. The following values are valid: integers between 0 and 100. (Superseded.)

Cumulative Process User CPU (Percent) The percentage of cumulative CPU time spent in user mode by process. The following values are valid: integers between 0 and 100. (Superseded.)

Data Resident Set (Pages) The size of the data set based on the number of pages. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Data Size (KB) The data size (in kilobytes) of the virtual memory. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Data Size (MB) The data size (in megabytes) of the virtual memory. This attribute displays as a floating point with a scale of 1. For example 5.2. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Executable Size (KB) The executable size (in kilobytes) of the virtual memory. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Executable Size (MB) The executable size (in megabytes) of the virtual memory. This attribute displays as a floating point with a scale of 1. For example 5.2. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Kernel Priority The kernel scheduling priority. The following values are valid: integers between -100 - 100 (-100 is the highest). Real-time processes can have priorities that are negative. (Superseded.)

Library Size (KB) The library size (in kilobytes) of the virtual memory. This measurement represents all pages, including unused. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Library Size (MB) The library size (in megabytes) of the virtual memory. This measurement represents all pages, including unused. This attribute displays as a floating point with a scale of 1. For example 5.2. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Nice Value The standard Linux nice level (-20 represents the highest level). The following values are valid: integers in the range -20 to 19. (Superseded.)

Number of Threads The number of threads started for this process. (Valid only on 2.6 kernel and above.) Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Parent Process ID The identifier for the parent process. The following values are valid: integers between 0 and 999. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Process Busy CPU (Percent) The summation of User CPU Percent and System CPU Percent for this process. (Superseded.)

Process Command Name The name of the process command. The following values are valid: alphanumeric text strings with a maximum length of 32 characters. (Superseded.)

Process Command Name (Unicode) The name of the process command. The following values are valid: text strings with a maximum length of 64 bytes. This attribute is globalized (Unicode). (Superseded.)

Process Count The count of processes with the same name. The name is selected by using the Command Line (UNICODE), CMDLINEU, attribute. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Process CPU ID The ID of the process CPU. The following values are valid: integers. Note: -1 indicates Not Available. (Superseded.)

Process Dirty Pages Pages that have been modified (dirty) in buffer (main memory), but not yet copied to the cache. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Process ID The identifier of the process. The following values are valid: integers between 0 and 999. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Process Short Term Avg Busy CPU (Percent) The summation of Proc System CPU Norm and Proc User CPU Norm for this process. CPU percentages are normalized to account for multiple online processors; percentages are normalized to a maximum of 100 percent. This metric is only available through situations and only when the Process name is also specified within the situation predicate. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Process Short Term Avg System CPU (Percent) The short term average of the percentage of CPU time spent in kernel mode by the process. CPU percentages are normalized to account for multiple online processors; percentages are normalized to a maximum of 100 percent. This metric is only available through situations and only when the Process name is also specified within the situation predicate. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Process Short Term Avg User CPU (Percent) The short term average of the percentage of CPU time spent in user mode by the process. CPU percentages are normalized to account for multiple online processors; percentages are normalized to a maximum of 100 percent. This metric is only available through situations and only when the Process name is also specified within the situation predicate. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Process State The state of the process (Sleeping, Disk, Running, Zombie, Trace, Dead, or N/A). (Superseded.) The following values are valid: integers between -1 and 5, where 0 = Sleeping, 1 = Disk, 2 = Running, 3 = Zombie, 4 = Trace, 5 = Dead, and -1 = Not_Available.

Process System CPU (Percent) The percentage of CPU time spent in kernel mode by process. The following values are valid: integers between 0 and 100. (Superseded.)

Process User CPU (Percent) The percentage of CPU time spent in user mode by process. The following values are valid: integers between 0 and 100. (Superseded.)

Resident Set Size (Pages) The number of pages the process has in real memory. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Session ID The session ID. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Shared Lib Resident Set (Pages) The number of pages of shared library set (mmap) memory. mmap is a system API that lets you map a file or device into memory. The mapped pages might be shared so that other processes can access them. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Shared Memory (Pages) The number of pages of shared (mmap) memory. mmap is a system API that lets you map a file or device into memory. The mapped pages might be shared so that other processes can access them. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Stack Size (KB) The stack size (in kilobytes) of the virtual memory. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. Note: -1 indicates `Not_Available` and -2 indicates `Not_Collected`. (Superseded.)

Stack Size (MB) The stack size (in megabytes) of the virtual memory. The following values are valid: integers. This attribute displays as a floating point with a scale of 1. For example 5.2. Note: the value -1 indicates `Not Available`, the value -2 indicates `Not Collected`, and the value 2147483647 indicates `Value_Exceeds_Maximum`. (Superseded.)

System Name The managed system name. The form should be `hostname:agent_code`. (Superseded.)

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

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.

Text Resident Set (Pages) The number of pages of text resident (mmap) memory. mmap is a system API that lets you map a file or device into memory. The mapped pages might be shared so that other processes can access them. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Total Major Faults The total number of major page faults (including child processes) since the start of the process. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Total Minor Faults The total number of minor page faults (including child processes) since the start of the process. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Total Size (Pages) The number of pages that the process has in real memory. The following values are valid: integers. (Superseded.) For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

User to System CPU (Percent) Of the total system CPU usage, the percentage that was user CPU usage. For example, 500% means that user CPU usage is 5 times the system CPU usage. The following values are valid: integers between -10,000 and 10,000. (Superseded.)

VM Locked Pages (KB) The size (in kilobytes) of locked pages of the virtual memory. The following values are valid: integers. Note: -1 indicates `Not_Available` and -2 indicates `Not_Collected`. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

VM Locked Pages (MB) The size (in megabytes) of locked pages of the virtual memory. This attribute displays as a floating point with a scale of 1. For example

5.2. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value -2147483648 indicates Value_Exceeds_Minimum. (Superseded.)

VM Size (KB) The size (in kilobytes) of the virtual memory. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

VM Size MB Virtual memory size in megabytes. This attribute displays as a floating point with a scale of 1. For example 5.2. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Process User Info attributes

The Process User Info attributes refer to characteristics associated with effective groups, file system groups, real groups, and saved groups.

Effective Group ID The identifier of the effective group. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Effective Group Name The effective group name. The following values are valid: text strings with a maximum length of 64 bytes.

Effective User ID The identifier of the effective user. The following values are valid: integers. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Effective User Name The name of the effective user. The following values are valid: text strings with a maximum length of 64 bytes.

File System Group Name The name of the file system group. The following values are valid: text strings with a maximum length of 64 bytes.

File System Group ID The identifier of the file system group. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. For example, the following value is valid:
Value_Exceeds_Maximum=2147483647.

File System User ID The identifier of the file system user. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

File System User Name The name of the file system user. The following values are valid: text strings with a maximum length of 64 bytes.

Process Command Line The Command Line string for the process.

Process Command Name Command name of the process.

Process ID The identifier associated with the process. The following values are valid: integers. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Process Parent ID The Parent Process ID. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Process State The state of the process (Sleeping, Disk, Running, Zombie, Trace, Dead, or N/A). The following values are valid: integers between -1 and 5, where 0 = Sleeping, 1 = Disk, 2 = Running, 3 = Zombie, 4 = Trace, 5 = Dead, and -1 = Not_Available.

Real Group ID The identifier of the real group. The following values are valid: simple text string, alphanumeric with a maximum length 16 characters. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Real Group Name The name of the real group. The following values are valid: simple text string, with a maximum length 64 bytes.

Real User ID The identifier of the real user. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Real User Name The name of the real user. The following values are valid: text strings with a maximum length of 64 bytes.

Saved Group ID The identifier of the saved group. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Saved Group Name The name of the saved group. The following values are valid: text strings with a maximum length of 64 bytes.

Saved User ID The identifier of the saved user. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Saved User Name The name of the saved user. The following values are valid: text strings with a maximum length of 64 bytes. This attribute is globalized (Unicode).

Session ID The session ID. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

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

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

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.

Terminal Device Name of the terminal device that started a process.

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

VM Size (MB) Virtual Memory Size in Megabytes. This attribute displays as a floating point with a scale of 1. For example 5.2. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Process User Info attributes (superseded)

The Process User Info attributes refer to characteristics associated with effective groups, file system groups, real groups, and saved groups. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Command Line (Unicode) Command Line string of the process. (Superseded.)

Effective Group ID The identifier of the effective group. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Effective Group Name The effective group name. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. (Superseded.)

Effective Group Name (Unicode) The effective group name. The following values are valid: text strings with a maximum length of 64 bytes. This attribute is globalized (Unicode). (Superseded.)

Effective User ID The identifier of the effective user. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Effective User Name The name of the effective user. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. (Superseded.)

Effective User Name (Unicode) The name of the effective user. The following values are valid: text strings with a maximum length of 64 bytes. This attribute is globalized (Unicode). (Superseded.)

File System Group Name The name of the file system group. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. (Superseded.)

File System Group Name (Unicode) The name of the file system group. The following values are valid: text strings with a maximum length of 64 bytes. This attribute is globalized (Unicode). (Superseded.)

File System Group ID The identifier of the file system group. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

File System User ID The identifier of the file system user. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

File System User Name The name of the file system user. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. (Superseded.)

File System User Name (Unicode) The name of the file system user. The following values are valid: text strings with a maximum length of 64 bytes. This attribute is globalized (Unicode). (Superseded.)

Process Command Name (Unicode) The Process Command name (Unicode). (Superseded.)

Process ID The identifier associated with the process. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Process Parent ID The Parent Process ID. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Process State The state of the process (Sleeping, Disk, Running, Zombie, Trace, Dead, or N/A). (Superseded.) The following values are valid: integers between -1 and 5, where 0 = Sleeping, 1 = Disk, 2 = Running, 3 = Zombie, 4 = Trace, 5 = Dead, and -1 = Not_Available.

Real Group ID The identifier of the real group. The following values are valid: simple text string, alphanumeric with a maximum length 16 characters. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Real Group Name The name of the real group. The following values are valid: simple text string, alphanumeric with a maximum length 16 characters. (Superseded.)

Real Group Name (Unicode) The name of the real group. The following values are valid: simple text string, with a maximum length 64 bytes. This attribute is globalized (Unicode). (Superseded.)

Real User ID The identifier of the real user. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Real User Name The name of the real user. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. (Superseded.)

Real User Name (Unicode) The name of the real user. The following values are valid: text strings with a maximum length of 64 bytes. This attribute is globalized (Unicode). (Superseded.)

Saved Group ID The identifier of the saved group. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Saved Group Name The name of the saved group. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. (Superseded.)

Saved Group Name (Unicode) The name of the saved group. The following values are valid: text strings with a maximum length of 64 bytes. (Superseded.)

Saved User ID The identifier of the saved user. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Saved User Name The name of the saved user. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. (Superseded.)

Saved User Name (Unicode) The name of the saved user. The following values are valid: text strings with a maximum length of 64 bytes. This attribute is globalized (Unicode). (Superseded.)

Session ID The session ID. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

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

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

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.

Terminal Device Name of the terminal device that started a process. (Superseded.)

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

VM Size (MB) Virtual Memory Size in Megabytes. This attribute displays as a floating point with a scale of 1. For example 5.2. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

RPC Statistics attributes

Use RPC Statistics to monitor remote procedure call (RPC) characteristics, such as the number of RPC server calls (including the number of rejected calls), packets that are not valid, and client calls.

RPC Calls Retransmitted The number of client calls that needed to be transmitted again. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

RPC Client Calls The number of calls to the server made by the clients of the server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

RPC Packets with Malformed Header The number of packets that were received at the server with header records that were not properly formatted. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

RPC Server Call Authorization Failures The number of packets that were received at the server with authorizations that were not valid. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

RPC Server Calls Rejected The number of calls made to the server, which were rejected. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

RPC Server Invalid Client Requests The number of packets that were received at the server, which had client requests that were not valid. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

RPC Total Server Calls Received The total number of calls made to the server (both valid and not valid). The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

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

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

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.

Times Authentication Refreshed The number of times the authentication of a client was refreshed. The following values are valid: integers. 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.

RPC Statistics attributes (superseded)

Use RPC Statistics to monitor remote procedure call (RPC) characteristics, such as the number of RPC server calls (including the number of rejected calls), packets that are not valid, and client calls. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

RPC Calls Retransmitted The number of client calls that needed to be transmitted again. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

RPC Client Calls The number of calls to the server made by the clients of the server. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

RPC Packets with Malformed Header The number of packets that were received at the server with header records that were not properly formatted. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

RPC Server Call Authorization Failures The number of packets that were received at the server with authorizations that were not valid. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

RPC Server Calls Rejected The number of calls made to the server, which were rejected. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

RPC Server Invalid Client Requests The number of packets that were received at the server, which had client requests that were not valid. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

RPC Total Server Calls Received The total number of calls made to the server (both valid and not valid). The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

Times Authentication Refreshed The number of times the authentication of a client was refreshed. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Sockets Detail attributes

The Sockets Detail attributes refer to characteristics associated with socket details, including user ID, local and foreign addresses, socket states, and socket protocols.

Foreign Address The address of the remote end of the socket. Like "netstat" * indicates that the address is unassigned/unavailable. The following values are valid: alphanumeric text strings with a maximum length of 16 characters.

Foreign Port The number of the foreign port. The following values are valid: integers. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Local Address The address of the local end of the socket, presented as a dotted ip address. The following values are valid: alphanumeric text strings with a maximum length of 16 characters.

Local Port The local port number. The following values are valid: integers. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Local Service Name The local port number translated to service name from /etc/services. The following values are valid: alphanumeric text strings with a maximum length of 64 characters.

Receive Queue (Bytes) The count of bytes not copied by the user program connected to this socket. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Send Queue (Bytes) The count of bytes not acknowledged by the remote host. The following values are valid: integers. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Socket Inode The inode used by the socket. The following values are valid: integers. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

Socket Owner Name The user name associated with the user ID that owns or started the socket connection. The following values are valid: text strings with a maximum length of 64 bytes.

Socket Protocol Indicates the sockets using this protocol. "Total" includes UNIX domain sockets not displayed here. The following values are valid: integers, where 0 = TCP, 1 = UDP, 2 = RAW, 3 = UNIX, -1 = Not Available, and -2 = Not Collected.

Socket State The state of the socket. The following values are valid: integers, where 1 = ESTABLISHED, 2 = SYN_SENT, 3 = SYN_RECV, 4 = FIN_WAIT1, 5 = FIN_WAIT2, 6 = TIME_WAIT, 7 = CLOSED, 8 = CLOSED_WAIT, 9 = LAST_ACK, 10 = LISTEN, 11 = CLOSING, and 12 = UNKNOWN.

Socket UID The user ID of the owner of the socket. The following values are valid: integers. For example, the following value is valid:
Value_Exceeds_Maximum=9223372036854775807.

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

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

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.

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

Sockets Detail attributes (superseded)

The Sockets Detail attributes refer to characteristics associated with socket details, including user ID, local and foreign addresses, socket states, and socket protocols. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Foreign Address The address of the remote end of the socket. Like "netstat" * indicates that the address is unassigned/unavailable. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. (Superseded.)

Foreign Port The number of the foreign port. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Local Address The address of the local end of the socket, presented as a dotted ip address. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. (Superseded.)

Local Port The local port number. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Local Service Name The local port number translated to service name from /etc/services. The following values are valid: alphanumeric text strings with a maximum length of 64 characters. (Superseded.)

Receive Queue (Bytes) The count of bytes not copied by the user program connected to this socket. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Send Queue (Bytes) The count of bytes not acknowledged by the remote host. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Socket Inode The inode used by the socket. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Socket Owner Name (Unicode) The user name associated with the user ID that owns or started the socket connection. The following values are valid: text strings with a maximum length of 64 bytes. This attribute is globalized (Unicode). (Superseded.)

Socket Protocol Indicates the sockets using this protocol. "Total" includes UNIX domain sockets not displayed here. The following values are valid: integers, where 0 = TCP, 1 = UDP, 2 = RAW, 3 = UNIX, -1 = Not Available, and -2 = Not Collected.

Socket State The state of the socket. (Superseded.) The following values are valid: integers, where 1 = ESTABLISHED, 2 = SYN_SENT, 3 = SYN_RECV, 4 = FIN_WAIT1, 5 = FIN_WAIT2, 6 = TIME_WAIT, 7 = CLOSED, 8 = CLOSED_WAIT, 9 = LAST_ACK, 10 = LISTEN, 11 = CLOSING, and 12 = UNKNOWN.

Socket UID The user ID of the owner of the socket. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

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

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

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Sockets Status attributes

The Sockets Status attributes refer to characteristics associated with the status of the Linux system sockets, including protocol names and high water marks used by protocols.

Highest Sockets Used The high water mark of sockets used by this protocol. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Socket Protocol Indicates the sockets using this protocol. "Total" includes UNIX domain sockets not displayed here. The following values are valid: integers, where 0 = TCP, 1 = UDP, 2 = RAW, 3 = UNIX, 4 = FRAG, -1 = TOTAL, and -2 = NOT_AVAILABLE.

Sockets in Use Sockets in use by protocol. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

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

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

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.

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

Sockets Status attributes (superseded)

The Sockets Status attributes refer to characteristics associated with the status of the Linux system sockets, including protocol names and high water marks used by protocols. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Highest Sockets Used The high water mark of sockets used by this protocol. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Socket Protocol Indicates the sockets using this protocol. "Total" includes UNIX domain sockets not displayed here. The following values are valid: integers, where 0 = TCP, 1 = UDP, 2 = RAW, 3 = UNIX, 4 = FRAG, -1 = TOTAL, and -2 = NOT_AVAILABLE.

Sockets in Use Sockets in use by protocol. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Swap Rate attributes

The Swap Rate attributes feature swap space characteristics, including usage rates and days till full data.

Days Until Swap Space Full The predicted number of days till swap space is completely used (moving average). The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647.

Low Water Mark for Free real memory (KB) The lowest level that Free real memory has reached, expressed in kilobytes. The following values are valid: integers. Note: -1 indicates Not Available and -2 indicates Not Collected. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Minimum Days to Swap Full The minimum number of days till swap space is completely used (peak rate based). The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647.

Peak Swap Space Used (MB) The peak swap space used based on snap shots, expressed in megabytes. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Swap Space Used (MB) (Moving Average) The moving average of swap space used, expressed in megabytes. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Swap Space Used (bytes per hour) The swap space usage rate, expressed in bytes per hour. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`.

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

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

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.

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

Total Swap Space (MB) (Moving Average) The moving average of total swap space, expressed in megabytes. The following values are valid: integers. For example, the following value is valid:
`Value_Exceeds_Maximum=9223372036854775807`.

Swap Rate attributes (superseded)

The Swap Rate attributes feature swap space characteristics, including usage rates and days till full data. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Days Until Swap Space Full The predicted number of days till swap space is completely used (moving average). The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Low Water Mark for Free real memory (KB) The lowest level that Free real memory has reached, expressed in kilobytes. The following values are valid: integers. Note: -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

Minimum Days to Swap Full The minimum number of days till swap space is completely used (peak rate based). The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Peak Swap Space Used (MB) The peak swap space used based on snap shots, expressed in megabytes. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Swap Space Used (MB) (Moving Average) The moving average of swap space used, expressed in megabytes. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Swap Space Used (bytes per hour) The swap space usage rate, expressed in bytes per hour. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Total Swap Space (MB) (Moving Average) The moving average of total swap space, expressed in megabytes. The following values are valid: integers. (Superseded.)

System Statistics attributes

The System Statistics attributes refer to characteristics associated with system performance such as the number of logged in users, the number of processes per second, and system load statistics.

Context Switches Per Second The number of context switches per second. Calculated on a 30 second interval. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Minimum=-9223372036854775808` and the value `Value_Exceeds_Maximum=9223372036854775807`.

Major Page Faults Per Second Number of major faults per second, these are page faults that directly require the loading of pages from disk. (Kernel 2.6 and greater.) Calculated on a 30 second interval. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value `9223372036854775807` indicates `Value_Exceeds_Maximum`.

Number of Processes in Zombie State Number of processes currently in Zombie State. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`.

Number of User Logins The current number of users logged in. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Page Faults Per Second The total number of page faults per second (major and minor). (Kernel 2.6 and above only.) Calculated on a 30 second interval. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value `9223372036854775807` indicates `Value_Exceeds_Maximum`.

Pages Paged In Per Second The pages paged in per second. Calculated on a 30 second interval. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`.

Pages Paged Out Per Second The pages paged out per second. Calculated on a 30 second interval. For example, the following value is valid: `Value_Exceeds_Maximum=9223372036854775807`.

Pages Swapped In The pages swapped in. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Pages Swapped In Per Second The pages swapped in per second. Calculated on a 30 second interval. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Pages Swapped Out The pages swapped out. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Pages Swapped Out Per Second The pages swapped out per second. Calculated on a 30 second interval. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Percent Change Context Switches Per Second The percentage change in the number of context switches per second. The following values are valid: integers in the range -100 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Percent Change Processes Created The percentage change in the number of processes created per second. The following values are valid: integers in the range -100 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Processes Created Per Second The number of processes created per second. Calculated on a 30 second interval. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. For example, the following value is valid: Value_Exceeds_Minimum=-9223372036854775808 and the value Value_Exceeds_Maximum=9223372036854775807.

System Load Last 1 Minute The load on the system for the last minute. Note that the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

System Load Last 5 Minutes The load on the system for the last 5 minutes. Note that the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

System Load Last 15 Minutes The load on the system for the last 15 minutes. Note that the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

System Name The managed system name. The form should be *hostname:agent_code*. Examples include spark:LZ or deux.raleigh.ibm.com:LZ.

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.

System Uptime The System Uptime in seconds, however it displays as a time counter on the Tivoli Enterprise Portal. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

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

Total Number of Processes The total number of processes. For example, the following value is valid: Value_Exceeds_Maximum=9223372036854775807.

Total Pages Paged In The total pages paged in. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Total Pages Paged Out The total pages paged out. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

System Statistics attributes (superseded)

The System Statistics attributes refer to characteristics associated with system performance such as the number of logged in users, the number of processes per second, and system load statistics. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Context Switches Per Second The number of context switches per second. Calculated on a 30 second interval. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Minimum=-2147483648 and the value Value_Exceeds_Maximum=2147483647. (Superseded.)

Major Page Faults Per Second Number of major faults per second, these are page faults that directly require the loading of pages from disk. (Kernel 2.6 and greater.) Calculated on a 30 second interval. (Superseded.) Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum.

Number of Processes in Zombie State Number of processes currently in Zombie State. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Number of User Logins The current number of users logged in. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Page Faults Per Second The total number of page faults per second (major and minor). (Kernel 2.6 and above only.) Calculated on a 30 second interval. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Pages Paged In Per Second The pages paged in per second. Calculated on a 30 second interval. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Pages Paged Out Per Second The pages paged out per second. Calculated on a 30 second interval. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Pages Swapped In The pages swapped in. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Pages Swapped In Per Second The pages swapped in per second. Calculated on a 30 second interval. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Pages Swapped Out The pages swapped out. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Pages Swapped Out Per Second The pages swapped out per second. Calculated on a 30 second interval. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Percent Change Context Switches Per Second The percentage change in the number of context switches per second. The following values are valid: integers in the range -100 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Percent Change Processes Created The percentage change in the number of processes created per second. The following values are valid: integers in the range -100 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

Processes Created Per Second The number of processes created per second. Calculated on a 30 second interval. The following values are valid: integers. Note: -1 indicates Not_Available and -2 indicates Not_Collected. (Superseded.)

System Load Last 1 Minute The load on the system for the last minute. The following values are valid: integers in the range 0 to 100. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

System Load Last 5 Minutes The load on the system for the last 5 minutes. The following values are valid: integers in the range 0 to 100. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

System Load Last 15 Minutes The load on the system for the last 15 minutes. The following values are valid: integers in the range 0 to 100. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

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

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

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.

System Uptime The System Uptime in seconds, however it displays as a time counter on the Tivoli Enterprise Portal. The following values are valid: integers. For example, the following value is valid: Value_Exceeds_Maximum=2147483647. (Superseded.)

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Total Number of Processes The total number of processes. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

Total Pages Paged In The total pages paged in. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates `Value_Exceeds_Maximum`. (Superseded.)

Total Pages Paged Out The total pages paged out. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates `Value_Exceeds_Maximum`. (Superseded.)

User Login attributes

The User Login attributes refer to user characteristics such as idle time, user name, location, and log in time.

Hostname (From) The host name associated with the log in for the user. The following values are valid: alphanumeric text strings with a maximum length of 256 characters.

Idle Time The number of minutes that have passed since a user last entered a command. The following values are valid: numeric values expressed as minutes in the range 0 to 20160. Use this attribute to check idle time.

Line The terminal device type or line to which the user is connected. The following values are valid: alphanumeric text strings with a maximum length of 16 characters.

Login Time The date and time the user logged in.

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

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

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.

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

User Login PID The log in ID of the user. The following values are valid: integers. For example, the following value is valid:
`Value_Exceeds_Maximum=9223372036854775807`.

User Name The full name of a user. The following values are valid: alphanumeric text strings with a maximum length of 96 characters.

User Login attributes (superseded)

The User Login attributes refer to user characteristics such as idle time, user name, location, and log in time. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Hostname (From) The host name associated with the log in for the user. The following values are valid: alphanumeric text strings with a maximum length of 256 characters. (Superseded.)

Idle Time The number of minutes that have passed since a user last entered a command. The following values are valid: numeric values expressed as minutes in the range 0 to 20160. Use this attribute to check idle time. (Superseded.)

Line The terminal device type or line to which the user is connected. The following values are valid: alphanumeric text strings with a maximum length of 16 characters. (Superseded.)

Login Time The date and time the user logged in. (Superseded.)

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

User Login PID The log in ID of the user. The following values are valid: integers. For example, the following value is valid: `Value_Exceeds_Maximum=2147483647`. (Superseded.)

User Name The full name of a user. The following values are valid: alphanumeric text strings with a maximum length of 32 characters. (Superseded.)

User Name (Unicode) The name of the user logging in to access the system. The following values are valid: text strings up to 64 bytes. This attribute is globalized (Unicode). (Superseded.)

VM Stats attributes

The VM Stats attributes refer to memory characteristics such as the size of cached, free, and shared memory.

Free Virtual Storage (MB) The available virtual storage (in megabytes). Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates `Value_Exceeds_Maximum`.

Free Virtual Storage (Percent) Available Virtual Storage in percentage. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Memory Cached (MB) The size (in megabytes) of physical memory cached. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates `Value_Exceeds_Maximum`.

Memory Cached Percent Physical memory cached, in percent. Note: the value -1 indicates Not Available and the value -2 indicates Not Collected.

Memory Free (MB) The size (in megabytes) of physical memory free. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Memory Free (Percent) The available real memory in percentage. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Memory in Buffers (MB) The size (in megabytes) of physical memory in buffers. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Memory Used (MB) The size (in megabytes) of physical memory used. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Memory Used (Percent) The used real memory in percentage. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Net Memory Used (MB) Physical Memory used (in MB) not considering Memory in Buffers and Memory Cached. It is calculated as Memory Used (MB) - Memory in Buffers (MB) - Memory Cached (MB). Note: the value -1 indicates Not Available, -2 indicates Not Collected, and 9223372036854775807 indicates Value_Exceeds_Maximum.

Net Memory Used (Percent) Physical Memory used (in Percent) not considering Memory in Buffers and Memory Cached. It is calculated as Memory Used (MB) - Memory in Buffers (MB) - Memory Cached (MB). Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Processes Blocked State The number of processes in uninterruptible sleep, waiting for I/O to complete. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Processes Waiting RunTime The number of processes waiting for CPU time in the run queue. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Shared Memory (MB) The size (in megabytes) of physical memory shared. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Swap Space Free (MB) The size (in megabytes) of swap space free. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Swap Space Free (Percent) Available Swap Space (Percent). Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Swap Space Used (MB) The size (in megabytes) of swap space used. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Swap Space Used (Percent) Used Swap Space (Percent). Note: the value -1 indicates Not Available and -2 indicates Not Collected.

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

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

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.

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

Total Memory (MB) The total size (in megabytes) of physical memory. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Total Memory Free (MB) Physical Memory free (in MB) considering Memory in Buffers and Memory Cached. It is calculated as Memory Free (MB) + Memory in Buffers (MB) + Memory Cached (MB). The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Total Memory Free (Percent) Physical Memory free (in percent) considering Memory in Buffers and Memory Cached. It is calculated as Memory Free (MB) + Memory in Buffers (MB) + Memory Cached (MB). The following values are valid: integers. Note: the value -1 indicates Not Available and the value -2 indicates Not Collected.

Total Swap Space (MB) The total size (in megabytes) of swap space. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Total Virtual Storage (MB) The total virtual storage (real plus swap storage) in MB. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Used Virtual Storage (MB) The used virtual storage in MB. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 9223372036854775807 indicates Value_Exceeds_Maximum.

Used Virtual Storage (Percent) The used virtual storage in percentage. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

VM Stats attributes (superseded)

The VM Stats attributes refer to memory characteristics such as the size of cached, free, and shared memory. This attribute group is superseded. There is a new attribute group with the same name that replaces it.

Available Virtual Storage (MB) The available virtual storage in MB. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Available Virtual Storage (Percent) The available virtual storage in percentage. Note: the value -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

Memory Cached (MB) The size (in megabytes) of physical memory cached. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Memory Free (MB) The size (in megabytes) of physical memory free. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Memory in Buffers (MB) The size (in megabytes) of physical memory in buffers. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Memory Used (MB) The size (in megabytes) of physical memory used. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Real Memory Available (Percent) Available Real Memory in Percent. Note: the value -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

Real Memory Used (Percent) Used Real Memory (Percent). Note: the value -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

Shared Memory (MB) The size (in megabytes) of physical memory shared. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Swap Space Available (Percent) Available Swap Space (Percent). Note: the value -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

Swap Space Free (MB) The size (in megabytes) of swap space free. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Swap Space Used (MB) The size (in megabytes) of swap space used. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Swap Space Used (Percent) Used Swap Space (Percent). Note: the value -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

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

Examples include *spark:LZ* or *deux.raleigh.ibm.com:LZ*.

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.

Time Stamp The date and time the agent collects information as set on the monitored system. (Superseded.)

Total Memory (MB) The total size (in megabytes) of physical memory. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Total Swap Space (MB) The total size (in megabytes) of swap space. The following values are valid: integers. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Total Virtual Storage (MB) The total virtual storage (real plus swap storage) in MB. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Used Virtual Storage (MB) The used virtual storage in MB. Note: the value -1 indicates Not Available, the value -2 indicates Not Collected, and the value 2147483647 indicates Value_Exceeds_Maximum. (Superseded.)

Used Virtual Storage (Percent) The used virtual storage in percentage. Note: the value -1 indicates Not Available and -2 indicates Not Collected. (Superseded.)

Disk capacity planning for historical data

Disk capacity planning for a monitoring agent is a prediction of the amount of disk space to be consumed for each attribute group whose historical data 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 Linux 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.

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)
KLZPASSTAT	KCA_LZ_Agent_Active_Runtime_Status	1410	1435	1472
KLZPASMGMT	KCA_LZ_Agent_Availability_Management_\\Status	554	559	596
KLZPASALRT	KCA_LZ_Alerts_Table	512	516	553
KLZPASCAP	KCA_LZ_Configuration_Information	3090	3132	3169
KLZCPU	KLZ_CPU	140	228	622
KLZCPUAVG	KLZ_CPU_Averages	160	313	1001
KLZDISK	KLZ_Disk	716	733	1098
KLZDSKIO	KLZ_Disk_IO	220	248	492
KLZDU	KLZ_Disk_Usage_Trends	436	447	812
KLZIOEXT	KLZ_IO_Ext	288	441	1294
KLZNFS	KLZ_NFS_Statistics	412	459	2107
KLZNET	KLZ_Network	385	411	1256
KLZPROC	KLZ_Process	1244	1455	3007
KLZPUSR	KLZ_Process_User_Info	1600	1632	1720
KLZRPC	KLZ_RPC_Statistics	172	176	429
KLZSOCKD	KLZ_Sockets_Detail	324	332	455
KLZSOCKS	KLZ_Sockets_Status	128	127	207
KLZSWPRT	KLZ_Swap_Rate	156	159	364
KLZSYS	KLZ_System_Statistics	264	361	1289
KLZLOGIN	KLZ_User_Login	516	522	559
KLZVM	KLZ_VM_Stats	228	338	1119

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)
LNXALLUSR	Linux_All_Users	180	182	219
LNXCPU	Linux_CPU	184	278	723
LNXCPUAVG	Linux_CPU_Averages	208	380	1170
LNXCPUCON	Linux_CPU_Config	328	335	372
LNXDISK	Linux_Disk	516	523	872
LNXDSKIO	Linux_Disk_IO	240	273	493
LNXDU	Linux_Disk_Usage_Trends	232	232	581
LNXFILCMP	Linux_File_Comparison	1652	1660	1697
LNXFILE	Linux_File_Information	3608	3653	3792
LNXFILPAT	Linux_File_Pattern	1652	1660	1697
LNXGROUP	Linux_Group	172	172	209
LNXPING	Linux_Host_Availability	244	255	343
LNXIOEXT	Linux_IO_Ext	276	474	1327
LNXPADDR	Linux_IP_Address	574	578	615
LNXMACHIN	Linux_Machine_Information	792	801	838
LNXNFS	Linux_NFS_Statistics	352	392	1740
LNXNET	Linux_Network	345	364	1025
LNXOSCON	Linux_OS_Config	468	460	497
LNXPROC	Linux_Process	1172	1393	2849
LNXPUSR	Linux_Process_User_Info	1432	1469	1557
LNXRPC	Linux_RPC_Statistics	180	177	334
LNXSCKD	Linux_Sockets_Detail	340	341	456
LNXSCKS	Linux_Sockets_Status	160	152	228
LNXSWPRT	Linux_Swap_Rate	176	172	365
LNXSYS	Linux_System_Statistics	232	350	1194
LNLOGIN	Linux_User_Login	552	552	589
LNVM	Linux_VM_Stats	220	371	1152

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

Note: The Linux Process attribute group is eligible for historical collection by default since the Linux Availability Historical workspaces require historical collection to be turned on for this attribute group. However, turning on historical collection for this attribute group is not recommended for all customers - customers who have large number of processes running on systems should weigh the costs (disk space, CPU, etc.) of collecting historical information on this attribute group.

Chapter 5. 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 shipped 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 Linux 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 by 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 *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 and the information in that section for each individual situation.

Predefined situations

This monitoring agent contains the following predefined situations:

The remaining sections of this chapter contain descriptions of each of these predefined situations. The situations are organized alphabetically.

- Linux_AMS_Alert_Critical

- Linux_BP_AvgCpuBusyPct1h_Critic
- Linux_BP_CpuBusyPct_Critical
- Linux_BP_CpuWaitIOPct_Warning
- Linux_BP_LoadAvg5min_Critical
- Linux_BP_NetTotalErrPct_Warning
- Linux_BP_NumberZombies_Warning
- Linux_BP_ProcHighCpu_Critical
- Linux_BP_ProcMissing_Critical
- Linux_BP_SpaceUsedPct_Critical
- Linux_BP_SpaceUsedPctCustom_Cri
- Linux_BP_SwapSpaceUsedPct_Criti
- Linux_Fragmented_File_System
- Linux_Fragmented_File_System_2
- Linux_High_CPU_Overload
- Linux_High_CPU_Overload_2
- Linux_High_CPU_System
- Linux_High_CPU_System_2
- Linux_High_Packet_Collisions
- Linux_High_Packet_Collisions_2
- Linux_High_RPC_Retransmit
- Linux_High_RPC_Retransmit_2
- Linux_High_Zombies
- Linux_High_Zombies_2
- Linux_Low_Pct_Inodes
- Linux_Low_Pct_Inodes_2
- Linux_Low_percent_space
- Linux_Low_percent_space_2
- Linux_Low_Space_Available
- Linux_Low_Space_Available_2
- Linux_Network_Status
- Linux_Network_Status_2
- Linux_NFS_Buffer_High
- Linux_NFS_Buffer_High_2
- Linux_NFS_Getattr_High
- Linux_NFS_Getattr_High_2
- Linux_NFS_rmlink_high
- Linux_NFS_rmlink_high_2
- Linux_NFS_Read_High
- Linux_NFS_Read_High_2
- Linux_NFS_Writes_High
- Linux_NFS_Writes_High_2
- Linux_Packets_Error
- Linux_Packets_Error_2
- Linux_Process_High_Cpu
- Linux_Process_High_Cpu_2

- Linux_Process_High_Instant_CPU
- Linux_Process_stopped
- Linux_Process_stopped_2
- Linux_RPC_Bad_Calls
- Linux_RPC_Bad_Calls_2
- Linux_System_Thrashing
- Linux_System_Thrashing_2

Linux_AMS_Alert_Critical situation

Monitors to determine 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'
```

Linux_BP_AvgCpuBusyPct1h_Critic situation

Monitors total percent 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.

The formula for this situation is as follows:

```
*IF *VALUE KLZ_CPU_Averages.CPU_Usage_Current_Average *GE 95.0
```

Linux_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 KLZ_CPU.CPU_ID *NE Aggregate *AND *VALUE KLZ_CPU.Busy
_CPU *GT 95.0 ) *OR ( *VALUE KLZ_CPU.CPU_ID *EQ Aggregate *AND *VALUE KL
Z_CPU.Busy_CPU *GT 90.0 ) )
```

Linux_BP_CpuWaitIOPct_Warning situation

Monitors CPU I/O waits. 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 KLZ_CPU.Wait_IO_CPU *GT 10 *AND *VALUE KLZ_CPU.CPU_ID *EQ  
Aggregate
```

Linux_BP_LoadAvg5min_Critical situation

Monitors the load on the system for 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 Linux_Machine_Information.Number_of_Processors_Online  
*EQ 1 *AND *VALUE KLZ_System_Statistics.System_Load_5min *GT 4.0 ) *OR  
( *VALUE Linux_Machine_Information.Number_of_Processors_Online *EQ 2 *AN  
D *VALUE KLZ_System_Statistics.System_Load_5min *GT 8 ) *OR ( *VALUE Lin  
ux_Machine_Information.Number_of_Processors_Online *EQ 3 *AND *VALUE KLZ  
_System_Statistics.System_Load_5min *GT 12 ) *OR ( *VALUE Linux_Machine_  
Information.Number_of_Processors_Online *EQ 4 *AND *VALUE KLZ_System_Sta  
tistics.System_Load_5min *GT 16 ) *OR ( *VALUE Linux_Machine_Information  
.Number_of_Processors_Online *EQ 6 *AND *VALUE KLZ_System_Statistics.Sys  
tem_Load_5min *GT 24 ) *OR ( *VALUE Linux_Machine_Information.Number_of_  
Processors_Online *EQ 8 *AND *VALUE KLZ_System_Statistics.System_Load_5m  
in *GT 32 ) )
```

Linux_BP_NetTotalErrPct_Warning situation

Monitors the percentage of errors of 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 KLZ_Network.Total_Error_Percent *GT 10
```

Linux_BP_NumberZombies_Warning situation

Monitors the number of processes for defunct. 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 KLZ_System_Statistics.Processes_Zombie_Count *GE 50
```

Linux_BP_ProcHighCpu_Critical situation

Monitors the CPU percent utilization by all processes except kswapd. 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 KLZ_Process.Busy_CPU_Pct *GE 95.0 *AND *VALUE KLZ_Process  
.Process_Command_Name *NE kswapd
```

Linux_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 KLZ_Process.Proc_CMD_Line *EQ ( '/usr/*/sshd', '/usr/*/s  
yslogd', '/usr/*/ntpd', '/usr/*/inetd', '/usr/*/crond' )
```

Linux_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 KLZ_Disk.Disk_Used_Percent *GE 95 *AND *VALUE KLZ_Disk.FS  
_Type *NE nfs
```

Linux_BP_SpaceUsedPctCustom_Cri 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 KLZ_Disk.Mount_Point *EQ '/' *AND *VALUE KLZ_Disk.Dis  
k_Used_Percent *GE 90 ) *OR ( *VALUE KLZ_Disk.Mount_Point *EQ '/home' *A  
ND *VALUE KLZ_Disk.Disk_Used_Percent *GE 80 ) ) *UNTIL ( *SIT Linux_BP_S  
paceUsedPct_Critical )
```

Linux_BP_SwapSpaceUsedPct_Criti 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 KLZ_VM_Stats.Swap_Used_Pct *GT 40
```

Linux_Fragmented_File_System situation

This situation has been superseded by Linux_Fragmented_File_System_2. Monitors the percentage of i-nodes to disk space. An exception condition occurs when the percentage of i-nodes to disk space used is high, which could indicate high disk fragmentation on the disk.

This situation has the following formula.

```
IF VALUE Linux_Disk.Space_Used_Percent LT 85 AND  
VALUE Linux_Disk.Inodes_Used_Percent GT 80
```

Linux_Fragmented_File_System_2 situation

Monitors the percentage of i-nodes to disk space. An exception condition occurs when the percentage of i-nodes to disk space used is high, which could indicate high disk fragmentation on the disk.

This situation has the following formula.

```
IF VALUE KLZ_Disk.Disk_Used_Percent LT 85 AND  
VALUE KLZ_Disk.Inodes_Used_Percent *GT 80
```

Linux_High_CPU_Overload situation

This situation has been superseded by Linux_High_CPU_Overload_2. Monitors the percentage of time the processor is busy. An exception condition occurs when the percentage is extremely high.

This situation has the following formula.

```
IF VALUE Linux_CPU.Idle_CPU LT 10.0 AND VALUE Linux_CPU.CPU_ID EQ Aggregate
```

Linux_High_CPU_Overload_2 situation

Monitors the percentage of time the processor is busy. An exception condition occurs when the percentage is extremely high.

This situation has the following formula.

```
IF VALUE KLZ_CPU.Idle_CPU LT 10.0 AND VALUE KLZ_CPU.CPU_ID EQ  
Aggregate
```

Linux_High_CPU_System situation

This situation has been superseded by Linux_High_CPU_System_2. Monitors the percentage of processor time that is used for system calls to check for runaway processes. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE Linux_CPU.CPU_ID EQ Aggregate AND VALUE  
Linux_CPU.System_CPU GT 80.0
```

Linux_High_CPU_System_2 situation

Monitors the percentage of processor time that is used for system calls to check for runaway processes. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE KLZ_CPU.CPU_ID EQ Aggregate AND VALUE KLZ_CPU.System_CPU  
GT 80.0
```

Linux_High_Packet_Collisions situation

This situation has been superseded by Linux_High_Packet_Collisions_2. Monitors the percentage of packet collisions during data transmission. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE Linux_Network.Collision_Percent GT 10
```

Linux_High_Packet_Collisions_2 situation

Monitors the percentage of packet collisions during data transmission. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE KLZ_Network.Collision_Percent GT 10
```

Linux_High_RPC_Retransmit situation

This situation has been superseded by Linux_High_RPC_Retransmit_2. Monitors the percentage of retransmits because of RPC Server calls. An exception condition occurs when the percentage is extremely high.

This situation has the following formula.

```
IF PCTCHANGE Linux_RPC_Statistics.RPC_Client_Calls_Retransmitted GT 10
```

Linux_High_RPC_Retransmit_2 situation

Monitors the percentage of retransmits because of RPC Server calls. An exception condition occurs when the percentage is extremely high.

This situation has the following formula.

```
IF PCTCHANGE KLZ_RPC_Statistics.RPC_Client_Calls_Retransmitted GT  
10
```

Linux_High_Zombies situation

This situation has been superseded by Linux_High_Zombies_2. Monitors the number of processes in zombie state. An exception condition occurs when the number is high.

This situation has the following formula.

```
IF VALUE Linux_Process.State EQ Zombie AND COUNT Linux_Process.State GT 20
```

Linux_High_Zombies_2 situation

Monitors the number of processes in zombie state. An exception condition occurs when the number is high.

This situation has the following formula.

```
IF VALUE KLZ_Process.State EQ Zombie AND COUNT  
KLZ_Process.Parent_Process_ID GT 20
```

Linux_Low_Pct_Inodes situation

This situation has been superseded by Linux_Low_Pct_Inodes_2. Monitors the percentage of available i-nodes. An exception condition occurs when the number is low.

This situation has the following formula.

```
IF VALUE Linux_Disk.Inodes_Used_Percent GT 80
```

Linux_Low_Pct_Inodes_2 situation

Monitors the percentage of available i-nodes. An exception condition occurs when the number is low.

This situation has the following formula.

```
IF VALUE KLZ_Disk.Inodes_Used_Percent GT 80
```

Linux_Low_percent_space situation

This situation has been superseded by Linux_Low_percent_space_2. Monitors the percentage of space available on a file system. An exception condition occurs when the percentage is low.

This situation has the following formula.

```
IF VALUE Linux_Disk.Space_Available_Percent LT 15
```

Linux_Low_percent_space_2 situation

Monitors the percentage of space available on a file system. An exception condition occurs when the percentage is low.

This situation has the following formula.

```
IF VALUE KLZ_Disk.Disk_Free_Percent LT 15
```

Linux_Low_Space_Available situation

This situation has been superseded by Linux_Low_Space_Available_2. Monitors the available space on a file system. An exception condition occurs when the amount of space is low.

This situation has the following formula.

```
IF VALUE Linux_Disk.Space_Available LT 7
```

Linux_Low_Space_Available_2 situation

Monitors the available space on a file system. An exception condition occurs when the amount of space is low.

This situation has the following formula.

```
IF VALUE KLZ_Disk.Disk_Free LT 7
```

Linux_Network_Status situation

This situation has been superseded by Linux_Network_Status_2. Monitors whether the Network Interface Card is up or not. An exception condition occurs when the network interface card is not up.

This situation has the following formula.

```
IF VALUE Linux_Network.Interface_Status NE UP
```

Linux_Network_Status_2 situation

Monitors whether the Network Interface Card is up or not. An exception condition occurs when the network interface card is not up.

This situation has the following formula.

```
IF VALUE KLZ_Network.Interface_Status NE UP
```

Linux_NFS_Buffer_High situation

This situation has been superseded by Linux_NFS_Buffer_High_2. Monitors the number of RPC retransmissions with no duplicate acknowledgements. An exception condition occurs when the number of retransmissions is high.

This situation has the following formula.

```
IF VALUE Linux_RPC_Statistics.RPC_Client_Calls_Retransmitted GT 60 AND  
PCTCHANGE Linux_RPC_Statistics.RPC_Client_Times_Authentication_Refreshed GT 5
```

Linux_NFS_Buffer_High_2 situation

Monitors the number of RPC retransmissions with no duplicate acknowledgements. An exception condition occurs when the number of retransmissions is high.

This situation has the following formula.

```
IF VALUE KLZ_RPC_Statistics.RPC_Client_Calls_Retransmitted GT 60  
AND PCTCHANGE KLZ_RPC_Statistics.RPC_Client_Times_Authentication_Refreshed  
GT 5
```

Linux_NFS_Getattr_High situation

This situation has been superseded by Linux_NFS_Getattr_High_2. Monitors the percentage of NFS server calls to read client attributes. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE Linux_NFS_Statistics.NFS_Get_Attribute_Calls_Pct GT 40
```

Linux_NFS_Getattr_High_2 situation

Monitors the percentage of NFS server calls to read client attributes. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE KLZ_NFS_Statistics.NFS_Get_Attribute_Calls_Pct GT 40
```

Linux_NFS_rmlink_high situation

This situation has been superseded by Linux_NFS_rmlink_high_2. Monitors the percentage of NFS server calls for read link operations. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE Linux_NFS_Statistics.NFS_Read_Link_Pct GT 10
```

Linux_NFS_rmlink_high_2 situation

Monitors the percentage of NFS server calls for read link operations. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE KLZ_NFS_Statistics.NFS_Read_Link_Pct GT 10
```

Linux_NFS_Read_High situation

This situation has been superseded by Linux_NFS_Read_High_2. Monitors the percentage of NFS server calls for read operations. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE Linux_NFS_Statistics.NFS_Read_Calls_Pct GT 30
```

Linux_NFS_Read_High_2 situation

Monitors the percentage of NFS server calls for read operations. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE KLZ_NFS_Statistics.NFS_Read_Calls_Pct GT 30
```

Linux_NFS_Writes_High situation

This situation has been superseded by Linux_NFS_Writes_High_2. Monitors the percentage of NFS server calls for write operations. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE Linux_NFS_Statistics.NFS_Writes_Pct GT 15
```

Linux_NFS_Writes_High_2 situation

Monitors the percentage of NFS server calls for write operations. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE KLZ_NFS_Statistics.NFS_Writes_Pct GT 15
```

Linux_Packets_Error situation

This situation has been superseded by Linux_Packets_Error_2. Monitors the percentage of network packets in error. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE Linux_Network.Total_Error_Percent GT 10
```

Linux_Packets_Error_2 situation

Monitors the percentage of network packets in error. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE KLZ_Network.Total_Error_Percent GT 10
```

Linux_Process_High_Cpu situation

This situation has been superseded by Linux_Process_High_Cpu_2. Monitors the percentage of processor time used by a process. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE Linux_Process.Busy_CPU_Pct GT 60.0
```

Linux_Process_High_Cpu_2 situation

Monitors the percentage of processor time used by a process. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE KLZ_Process.Busy_CPU_Pct GT 60.0
```

Linux_Process_High_Instant_CPU situation

Monitors the CPU usage percentage used by a process. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF VALUE KLZ_Process.Instant_Busy_CPU_Pct GT 60.0
```

Linux_Process_stopped situation

This situation has been superseded by Linux_Process_stopped_2. Monitors the number of stopped processes on the system. An exception condition occurs when the number is high.

This situation has the following formula.

```
IF VALUE Linux_Process.State NE Running AND  
VALUE Linux_Process.State NE Sleeping
```

Linux_Process_stopped_2 situation

Monitors the number of stopped processes on the system. An exception condition occurs when the number is high.

This situation has the following formula.

```
IF VALUE KLZ_Process.State NE Running AND VALUE  
KLZ_Process.State NE Sleeping
```

Linux_RPC_Bad_Calls situation

This situation has been superseded by Linux_RPC_Bad_Calls_2. Monitors the percentage of rejected RPC server or client calls. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF ( ( VALUE Linux_RPC_Statistics.RPC_Client_Calls_Retransmitted GT 30
) OR ( VALUE Linux_RPC_Statistics.RPC_Server_Calls_Rejected GT 30 ) )
```

Linux_RPC_Bad_Calls_2 situation

Monitors the percentage of rejected RPC server or client calls. An exception condition occurs when the percentage is high.

This situation has the following formula.

```
IF ( ( VALUE KLZ_RPC_Statistics.RPC_Client_Calls_Retransmitted
GT 30 ) OR ( VALUE KLZ_RPC_Statistics.RPC_Server_Calls_Rejected GT 30 ) )
```

Linux_System_Thrashing situation

This situation has been superseded by Linux_System_Thrashing_2. Monitors the swap space paging activity on the system. An exception condition occurs when the activity is extremely high.

This situation has the following formula.

```
IF ( ( VALUE Linux_System_Statistics.Pages_paged_out_per_sec GT 400.0 )
OR ( *VALUE Linux_System_Statistics.Pages_paged_in_per_sec GT 400.0 ) )
```

Linux_System_Thrashing_2 situation

Monitors the swap space paging activity on the system. An exception condition occurs when the activity is extremely high.

This situation has the following formula.

```
IF ( ( VALUE KLZ_System_Statistics.Pages_paged_out_per_sec GT 400
.0 ) OR ( VALUE KLZ_System_Statistics.Pages_paged_in_per_sec
GT 400.0 ) )
```

Chapter 6. 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 executes 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.

For more information about working with Take Action commands, see the *IBM Tivoli Monitoring User's Guide*.

Predefined Take Action commands

This monitoring agent offers 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 following information is provided about the Take Action command:

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 executed: 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 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 Linux OS with this action. Only the other agents being managed by Agent Management Services running on the Monitoring Agent for Linux 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 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 Linux OS with this action. Only the other agents being managed by Agent Management Services running on the Monitoring Agent for Linux 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 Linux OS with this action. Only the other agents being managed by Agent Management Services running on the Monitoring Agent for Linux 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 is what provides auto-restart capability.

Destination systems

Managed system

Usage notes

You now can target the Monitoring Agent for Linux 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 Linux 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 Start Manage action will not be allowed against any other non-OS agent. The `Linux_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 Linux OS will execute 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 executed directly by the remote Monitoring Agent for Linux 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

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 executes 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: For monitoring agents that provide predefined policies, predefined policies 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.

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 and the information in that section for each individual policy.

Predefined policies

There are no predefined policies for this monitoring agent.

Chapter 8. Server dashboards

The IBM Tivoli Monitoring Infrastructure Management Dashboards for Servers is a web-based application that runs in the Dashboard Application Services Hub. The server dashboards give the overall status of the service areas in your managed network.

Server dashboards background information

Use the Infrastructure Management Dashboards for Servers to assess the event and system status of your managed network that is filtered by your area of responsibility.

The information in the server dashboards ranges from a high-level overview of all managed system groups, to more detailed dashboards with key performance information about the selected group, managed system, or situation event.

Server dashboards and the Tivoli Enterprise Portal

After your administrator configures a connection to the dashboard data provider in the Dashboard Application Services Hub console, you can log in to the console from your browser and view the server dashboards. Situation event status information and operating system metrics that are obtained from your Tivoli Monitoring environment are displayed in the server dashboards.

Use the Tivoli Enterprise Portal or the command-line interface to complete tasks such as situation editing and historical reporting. You can launch the Tivoli Enterprise Portal from the server dashboards.

For the procedure to connect to the dashboard data provider, see “Preparing your dashboard environment” in the *IBM Tivoli Monitoring Administrator’s Guide*.

For a list of supported browsers, see “Supported browser versions” in the *IBM Tivoli Monitoring Installation and Setup Guide*.


For details about the Tivoli Enterprise Portal, see the *Tivoli Enterprise Portal User’s Guide* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/adminuse/itm_tepuser.htm)

For details about the tacmd CLI, see the *IBM Tivoli Monitoring Command Reference* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/cmdref/itm_cmdref.htm).

Managed system groups

Managed system groups are named lists of systems where Tivoli Enterprise Monitoring Agents are installed and are typically organized by function or line of business. Every Tivoli Monitoring product has at least one predefined managed system group, indicated by an asterisk at the beginning of the list name, such as *NT_SYSTEM for the Windows OS agent. You can create custom managed system groups in the Tivoli Enterprise Portal client or the command-line interface. (See “Managed system groups” in the *Tivoli Enterprise Portal User’s Guide*.)

Server dashboards



In the Dashboard Application Services Hub console, click  **System Status and Health** > **Server Dashboards** to open the home dashboard.

Managed System Groups Overview shows the situation event status of all the managed system groups in your managed network. From there, you can link to a tabbed dashboard with event information about the chosen managed system group, and link to detailed metrics for an individual managed system or situation event. You can also link to the detailed metrics from the **Situation Events** dashboard.

The server dashboards are all predefined and cannot be edited. You can, however, change the time span reported in the situation event results and adjust the display of tables and charts.

Note that in bi-directional language locales, the charts do not have full bi-directional support: the chart elements are not mirrored, but the legend and labels are.

Dashboard navigation

Along the top of every dashboard, is the path to the current dashboard. Click one of the hypertext links in the path to return to a previous dashboard. On the left side of the page are resource navigation icons ( and ) that you can click to return to the home dashboard or to open the Situation Events dashboard.

Managed System Groups Overview dashboard

This is the home dashboard gives a high-level overview of the open situation events and their severity for every managed system group associated with the hub Tivoli Enterprise Monitoring Server.

The top view shows the managed system groups as icons. Click **Switch View** to toggle between the icon view and a table view of the situations events that are open for each managed system group. From either view, you can link to a dashboard with metrics from the selected managed system group. From the table view, you can also link to specific event details or to key metrics on the managed system.

The bar charts along the bottom of the dashboard show a total count of the situation events in the managed network, one bar for each severity; the most critical managed system groups, one bar for each group and ordered by those with the most severity; and the situation event count for each monitoring agent type. For more information, see “Managed System Groups Overview dashboard” on page 132.

Managed system group dashboard

From the Managed System Groups Overview dashboard, select a managed system group such as *LINUX_SYSTEM to link to a tabbed page of situation status and key operating system metrics for each of the managed systems in the group. For more information, see “Managed system group dashboard” on page 133.

Managed system dashboard

From the managed system group dashboard Overview tab and from the Situation Events tables, you can click a managed system name to open a dashboard of detailed metrics and event listing for the selected managed system. For more information, see “Linux managed system dashboard” on page 138.

Situation Events dashboard



The Situation Events dashboard provides an alternate overview. Here you see a table of the situation event status from all the managed systems that report to the hub monitoring server. For more information, see “Situation Events dashboard” on page 135.

You can see a filtered version of the same table in the Situation Events tab of the managed system group dashboard and the managed system dashboard. For more information, see “Displaying situation event results” on page 131.

Situation event results dashboard

The event results dashboard gives you a chart for every expression in the formula, plotted with two hours of data samples before the event and up to two hours after event. You can extend this range by as much time as you have historical data for, and read the charts to help identify trends and data spikes over time. The dashboard is available by clicking a situation event from the Situation Events table. For more information, see “Situation event results dashboard” on page 136.

Returning to a previously viewed dashboard or exiting the web application

Use the resource navigation icons ( and ) to go back to a previous dashboard rather than your browser's Back button. Dashboards are treated as a single portal page; thus, using the browser's Back button exits the application.

After you exit the console with the logout option, by closing the browser window, or by clicking the browser's Close button or Back button, it might take a moment for proper shutdown. Do not repeatedly click the Close or Back button or the browser might become unresponsive.

Tip: The Infrastructure Management Dashboards for Servers are designed for a screen resolution of 1280 x 1024 pixels. For the best fit of dashboard elements, adjust your browser window to 1280 x 1024 pixels, or maximize the window if your monitor is at that resolution.

Role-based authorization policy


Tivoli Enterprise Portal permissions is the default authorization type. When this type of authorization is used, the administrator assigns monitoring applications and, optionally, permission to view events to your Tivoli Enterprise Portal user ID. Such permissions give you access to all managed systems and managed system groups for the monitoring application, for example, access to all Windows OS agent managed systems. Alternatively, the administrator can use authorization policies to give you access to specific managed system groups and managed systems, for example, access to a subset of the Windows OS agent managed systems and managed system groups.

When your Dashboard Application Services Hub configuration includes the Authorization Policy Server, the administrator can control user access to managed systems groups and individual managed systems. It is possible for the roles that are assigned to your user ID to have a mix of permissions that make it possible, for example, to see the dashboard for a particular managed system but not the situation events for that managed system type. In this case, the dashboard is displayed but not the situation event data.

If you open the server dashboards and get empty charts and tables but no error messages, it might be because you are not authorized to see any data.

Supported monitoring agents

In this release of Tivoli Monitoring, the supported monitoring agents are the OS agents for Linux, UNIX, and Windows.


If the managed system group that you are viewing includes managed systems that are not supported, the situation event results and metrics for those managed systems are excluded from most of the consolidated views of the managed system group dashboards and the situation events dashboards. However, they are included in the situation event counts, such as 7. If you open the situation event results dashboard for an unsupported monitoring agent, the Details and Advice tabs are empty except for the formula.

Actions

All server dashboards have an **Actions** menu from which you can select options to launch into the Tivoli Enterprise Portal, copy the URL of the current dashboard, set log tracing levels for diagnosing problems, and for displaying the server dashboards version information. For more information, see “Copying the URL” on page 142, “Launching to the Tivoli Enterprise Portal” on page 142, and “Setting a trace” on page 143.

Dashboard Health Checks

You can open the Dashboard Health Checks to run a health check of your infrastructure management dashboards components and report their status:

Click  **System Status and Health > Dashboard Health Checks**. The connection to the dashboard data provider on the Tivoli Enterprise Portal Server is checked and the results are shown in the Tivoli Monitoring table.

The Virtual Infrastructure Monitoring table shows the results of similar health checks for the installed dashboards, such as the Dashboard for VMware. If the table is empty, the IBM Infrastructure Management Dashboard for VMware is not installed on the Dashboard Application Services Hub.

See also the “Infrastructure Management Dashboards troubleshooting” topics in the *IBM Tivoli Monitoring Troubleshooting Guide*.

Related reference:

“Managed System Groups Overview dashboard” on page 132

Open the **Managed System Groups Overview** dashboard to see the consolidated situation event status for each managed system group in your environment.

“Situation Events dashboard” on page 135

Use the Situation Events dashboard for an overview of all the open events in your managed network.


Checking the health of your monitored environment

Open the Infrastructure Management Dashboards for Servers to see all or a subset of the managed system groups that you are responsible for in your managed enterprise. Looking at the situation event status and severity, you can quickly see problem areas.

Procedure

1. If you are not already logged on to the Dashboard Application Services Hub, log on now with the URL and user ID provided by your administrator. The

URL is `http://computer_name:16310/ibm/console` or `https://computer_name:16311/ibm/console` where *computer_name* is the fully qualified host name or IP address of the computer where the Dashboard Application Services Hub is installed. If your environment was configured with a port number other than the default, enter that number instead. The default path to the server is `/ibm/console`. However, this path is configurable, and might differ from the default in your environment.

2. In the navigation bar, select  **System Status and Health > Server Dashboards**.

Results

The Managed System Groups Overview dashboard is displayed with situation event metrics for all the managed system groups that are associated with the hub Tivoli Enterprise Monitoring Server.

What to do next

Click the **Learn more** link in the Managed System Groups Overview dashboard or one of the “Related reference” links to learn about the metrics that are displayed and what you can do in the dashboard.

Related concepts:

“Server dashboards background information” on page 127

Use the Infrastructure Management Dashboards for Servers to assess the event and system status of your managed network that is filtered by your area of responsibility.

Related reference:

“Managed System Groups Overview dashboard” on page 132

Open the **Managed System Groups Overview** dashboard to see the consolidated situation event status for each managed system group in your environment.



“Situation Events dashboard” on page 135

Use the Situation Events dashboard for an overview of all the open events in your managed network.

Displaying situation event results

Open the event dashboard for a situation to see the value or values that triggered the event. Also shown is a range of data samples that were taken before and after the event.

Procedure

1. Open the Situation Events table:
 - To see all events in the managed network, click  **Situation Events**.
 - To see all events in a managed system group, click  **Managed System Groups**, click the managed system group link from the carousel or scorecard view, and click the **Situations** tab.
 - To see all events on a single managed system, open the managed system dashboard from one of three places:
 - In the Situation Events dashboard, click the link in the Source column.
 - In the Managed System Group dashboard, click the managed system link from the carousel or scorecard view of the **Overview** tab.

- In the Managed System Group dashboard, click the link in the Source column of the **Situation Events** tab.
2. In the Situation Events table, click the link in the Situation Name column for the event results to display.

Results

The situation event results dashboard, which is named for the situation, such as “Linux Disk Critical”, is displayed with situation event metrics to help you determine the cause of the event.

What to do next

You can turn off the display of the thresholds broken line in the charts, lengthen the time range that is shown before and after the event, and control which event metrics are displayed in the chart. For more information, see “Situation event results dashboard” on page 136.

You can also launch into the Tivoli Enterprise Portal for closer scrutiny and to take further action. For more information, see “Launching to the Tivoli Enterprise Portal” on page 142.

Managed System Groups Overview dashboard

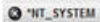
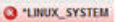
Open the **Managed System Groups Overview** dashboard to see the consolidated situation event status for each managed system group in your environment.

The event count by severity is displayed for every managed system group in your managed network, and filtered by what is allowed for your user profile.

The following views are displayed in the Managed System Groups Overview dashboard:

Managed System Groups

A count of the situation events by severity is shown for each managed system group in your managed enterprise, sorted by the groups with the most high severity situations. The following examples show how two managed system groups are ordered by severity in the carousel and scorecard views (user clicks **Switch View**).

Click the hypertext link of a managed system group, such as  or , to open a dashboard of situation event metrics for that managed system group (see “Managed system group dashboard” on page 133).



Situation Event Count by Severity

The total number of Fatal, Critical, Warning, Harmless, Informational, and Unknown events in the managed network is plotted, with one bar for each severity.

Situation Event Count by Managed System Group

The total number of open events is plotted, one bar for each managed system group. The event count in each managed system is shown in a stacked bar, with a slice for each event severity.

Situation Event Count by Managed System Type

The total number of open events is plotted, one bar for each managed system type in the managed network. The event count by each managed system type is shown in a stacked bar, with a slice for each event severity.

Related concepts:

“Server dashboards background information” on page 127

Use the Infrastructure Management Dashboards for Servers to assess the event and system status of your managed network that is filtered by your area of responsibility.

Related reference:

“Page layout and controls” on page 140

Use the server dashboard tools for manipulating the page and accessing other dashboards.

“Situation Events dashboard” on page 135

Use the Situation Events dashboard for an overview of all the open events in your managed network.

“Managed system group dashboard”


Use the *managed system group* dashboard to get more information about the selected managed system group.

“Linux managed system dashboard” on page 138

Use the Linux managed system dashboard to get a status overview of situation events and to see system details.

Managed system group dashboard

Use the *managed system group* dashboard to get more information about the selected managed system group.

This dashboard is opened from the Managed System Groups Overview dashboard by clicking the group name hypertext link such as  from either the carousel or scorecard view of the situation events.

Overview tab

The **Overview** tab is displayed with the metrics about the situation events opened for the managed system group.

Managed system group situation events by host name

A count of the situations events by severity is shown for each managed system in the managed system group, sorted by the groups with the most high severity situations.

- Click the hypertext link of a managed system to drill down to the situation event metrics for that system and to see key performance indicators for the operating system.

- Click the **Switch View** button to switch between the carousel view and scorecard view of the managed systems. The table view gives the same count of events by severity as the icon view, one row per managed system.

Situation Event Count by Severity

The total number of Fatal, Critical, Warning, Harmless, Informational, and Unknown events in the managed system group, one bar for each severity.

Most Critical Servers














The total number of open events for each managed system is plotted. The event count in each managed system is shown in a stacked bar, with a slice for each event severity.

Situation Event Count by Managed System Type

The total number of open events is plotted, one bar for each managed system type in the managed system group. The event count by each managed system type is shown in a stacked bar, with a slice for each event severity.

Situation Events tab

The Situation Events tab shows a table with all the events for the managed system group.

- For each managed system, the situation events and their status is displayed, sorted by the events with the highest severity first, and refreshed as new events arrive.
- You can click  **Pause updates** to temporarily stop automatic refresh as new events are opened; click  to resume.
- The toolbar shows a count of events for each severity. Click one or more of the tools to filter the list by event severity or status. For example, you can filter the list to show only open events with a severity of fatal or critical.
 -  Show fatal events
 -  Show critical events
 -  Show warning events
 -  Show harmless events
 -  Show informational events
 -  Show unknown events
 -  Show open status
 -  Show acknowledged status
 -  Show stopped status
 -  Show problem status
 -  Show expired status (acknowledgement expired and the situation is still true)

The tools toggle the filters off and on: Click a tool again to remove the filter

- Use the filter field to locate a situation by its name, display item, type or timestamp. See also “Table controls” on page 141.
- Click a hypertext link in the **Situation Name** column to open a dashboard of event details and expert advice. See “Displaying situation event results” on page 131 for a description of the event dashboard.
- Click a hypertext link in the **Source** column to open a dashboard of key performance metrics from the managed system and a table of situation events on the managed system. See “Linux managed system dashboard” on page 138.

Related reference:

“Page layout and controls” on page 140

Use the server dashboard tools for manipulating the page and accessing other dashboards.














“Managed System Groups Overview dashboard” on page 132

Open the **Managed System Groups Overview** dashboard to see the consolidated situation event status for each managed system group in your environment.

Situation Events dashboard

Use the Situation Events dashboard for an overview of all the open events in your managed network.

You can change the sort order and filter the table by event status or by cell values, and drill down to event details or to operating system details about the managed system.

- For every managed system group, the Situation Events dashboard displays the situation events that were opened and their status. The events are sorted by highest severity first and the dashboard is refreshed as new events arrive.
- You can click  **Pause updates** to temporarily stop automatic refresh as new events are opened; click  to resume.
- The toolbar shows a count of events for each severity. Click one or more of the tools to filter the list by event severity or status. For example, you can filter the list to show only open events with a severity of fatal or critical.
 -  Show fatal events
 -  Show critical events
 -  Show warning events
 -  Show harmless events
 -  Show informational events
 -  Show unknown events
 -  Show open status
 -  Show acknowledged status
 -  Show stopped status
 -  Show problem status
 -  Show expired status (acknowledgement expired and the situation is still true)

The tools toggle the filters off and on: Click a tool again to remove the filter

- Use the filter field to locate a situation by its name, display item, type or timestamp. See also “Table controls” on page 141.
- Click a hypertext link in the **Situation Name** column to open a dashboard of event details and expert advice. See “Displaying situation event results” on page 131 for a description of the event dashboard.
- Click a hypertext link in the **Source** column to open a dashboard of key performance metrics from the managed system and a table of situation events on the managed system. See “Linux managed system dashboard” on page 138.

Related concepts:

“Server dashboards background information” on page 127

Use the Infrastructure Management Dashboards for Servers to assess the event and system status of your managed network that is filtered by your area of responsibility.

Related reference:

“Page layout and controls” on page 140

Use the server dashboard tools for manipulating the page and accessing other dashboards.

“Managed System Groups Overview dashboard” on page 132

Open the **Managed System Groups Overview** dashboard to see the consolidated situation event status for each managed system group in your environment.

Situation event results dashboard

Open the event dashboard for a situation to see the value or values that triggered the event and a range of data samples before and after the event.


You can link to the dashboard by clicking the hypertext link in the **Situation Name** column of the **Situation Events** table, which is displayed in the Situation Events dashboard, the Situation Events tab of the managed system group dashboard, and the Overview tab of the managed system dashboard.

Details tab

Line chart

One line chart is displayed for each numeric attribute in the situation formula. The data sample, or samples if the situation has multiple expressions, that triggered the event is shown, along with up to two hours of historical data before and after the event.

The solid line plots the values of the attribute that triggered the event. The gray area begins at the situation interval when the threshold was breached and the event was opened. In the following example, the situation is true if the free disk space falls below 5%. We can see that the value falls below 5% a few minutes before the event is opened. The event occurrence takes place after the threshold was exceeded rather than at the same time because the historical data sampling is recorded more frequently than the sampled situation interval, which, in this case, is 15 minutes. The gray area continues on the time line while the condition remains true.



The broken line along the X-axis marks the threshold that was set. If you have multiple expressions for the same attribute, multiple threshold lines are displayed. Clear the **Show Thresholds** check box if you want to remove the threshold lines from the chart.

Viewing a range of data samples helps you to identify trending behavior. You can use the time selector to show longer periods of time before the event. Click the time selector bar (such as **Last 2 Hour(s) Before Event Occurrence**) and click one of the options to increase the time that is shown before the event to 4, 8, or 12 hours; 1 day; 1 week; 1, 3, 6, 9, or 12 months; or to customize the time range. For details, see “Select Time Range.”

Note: No historical data samples are plotted if historical collection was not configured for the attribute group or if the collection was only recently configured. For example, you must wait at least two hours after you configure a new historical collection before you can expect to see historical values if data samples are saved only once per hour. For more information about historical data collection, see “Creating a historical collection” in the *Tivoli Enterprise Portal User’s Guide* or the tacmd commands for history in the *IBM Tivoli Monitoring Command Reference* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/cmdref/itm_cmdref.htm).

Text values (such as missing processes) cannot be plotted in a chart.

Table Below the formula is a table with the data sample that triggered the event. For events that include a current data sample, you have an Expand icon in the **Server Name** column that you can click to see the current results. If the table has multiple rows, you can filter or move through the table pages as described in “Table controls” on page 141.

The first column shows the color that corresponds to the plot points for that attribute in the line chart. You can clear the check box to remove the line from the chart.

Advice tab

The **Advice** tab shows any expert advice that was written for the situation. If you created a situation in the Tivoli Enterprise Portal or Tivoli Monitoring command-line interface with advice that includes a link to a file, the URL is displayed. For more information, see “Writing expert advice” in the *Tivoli Enterprise Portal User’s Guide* and the tacmd commands for situations in the *IBM Tivoli Monitoring Command Reference* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/cmdref/itm_cmdref.htm).

Select Time Range




After you select **Custom** from the time selector options for the situation event results charts, the Select Time Range dialog box opens for you to specify a time range of values to see before or after (or both) the event time.

Select one of the following options:

- **Last**

Type or select the number and select the time unit for the time period before the event occurred: Minutes, Hours, Days, or Weeks. A default post-event time period of up to two hours is also plotted.

- **Set Time Range**

Drag the Start and End selectors (click  or  to move another time range into focus) to change Start and End Date and Time fields. You can also type directly in the fields or click the  list to see and select from a calendar and time range. The time range selection must include the event date and time.

After you click **OK**, the specified time range is shown in the event results dashboard. If you close the dashboard and reopen it, the time range returns to the two hours surrounding the event occurrence, which is the default setting.

Linux managed system dashboard

Use the Linux managed system dashboard to get a status overview of situation events and to see system details.

After you select the Linux OS managed system link from the Managed System Group dashboard or from the Situation Events dashboard's Source column, the Linux managed system dashboard is displayed.

The dashboard has an **Overview** tab where you can see overview charts of key performance indicators and a Situations Events table of open events for the selected managed system, and a **Process** tab showing current process utilization information in bar charts and a Process Table. (See also "Page layout and controls" on page 140.)

Overview tab

CPU Utilization (%) - Top 5

Shown here is a bar for each of the five processes that use the highest percentage of CPU on the selected managed system. If CPU usage is high, it is likely that one or more of these processes is responsible.

Memory Utilization (%) - Top 5

This bar chart shows the five processes that use the highest percentage of memory on the selected managed system.

Disk Utilization (%)

Shown here is a bar for each of the five processes that consume the highest percentage of disk space compared with the space allocated to them. If disk allocation is high, consider increasing the allocated space.



Situation Event Count by Severity

The total number of Fatal, Critical, Warning, Harmless, Informational, and Unknown events that are opened for the managed system, one bar for each severity.












Network Utilization (Packets/Second) - Top 5

Shown here are the five network interface adapters that send and receive the most packets per second. From this stacked bar chart, you can see where the most traffic occurs and how efficient the load balancing is.

Situation Events

- The situation events for the managed system and their status are displayed, sorted by the events with the highest severity first, and refreshed as new events arrive.
- You can click  **Pause updates** to temporarily stop automatic refresh as new events are opened; click  to resume.

- The toolbar shows a count of events for each severity. Click one or more of the tools to filter the list by event severity or status. For example, you can filter the list to show only open events with a severity of fatal or critical.

-  Show fatal events
-  Show critical events
-  Show warning events
-  Show harmless events
-  Show informational events
-  Show unknown events
-  Show open status
-  Show acknowledged status
-  Show stopped status
-  Show problem status
-  Show expired status (acknowledgement expired and the situation is still true)

The tools toggle the filters off and on: Click a tool again to remove the filter

- Use the filter field to locate a situation by its name, display item, type or timestamp. See also “Table controls” on page 141.
- Click a hypertext link in the **Situation Name** column to open a dashboard of event details and expert advice. See “Displaying situation event results” on page 131 for a description of the event dashboard.

Process tab

The **Process** tab gives key values related to the Linux processes running on the managed system.

Process CPU (%) Utilization - Top 5

Shown here are the five processes that use the most CPU compared with the space allocated to them. If process allocation is high, consider increasing the allocated space.

Process + Child CPU (%) Utilization - Top 5

Shown here are the five processes and child processes that use the most CPU compared with the space allocated to them. If process allocation is high, consider increasing the allocated space.

CPU Utilization (%) - Top 5

Shown here is a bar for each of the five processes that use the highest percentage of CPU on the selected managed system. If CPU usage is high, it is likely that one or more of these processes is responsible.

Memory Utilization (%) - Top 5

This bar chart shows the five processes that use the highest percentage of memory on the selected managed system.

Process Table

The Process Table shows process metrics for the selected managed system, one row per process, with the number of megabytes used compared with what was allocated, expressed as CPU percentage. The CPU percentage columns (System CPU, User CPU, Cumulative System CPU, and

Cumulative User CPU) provide a linear gauge for quickly identifying high CPU rates. The Command Line column shows the path and command that was issued to start the process.

See also “Table controls” on page 141.

Related reference:

“Managed System Groups Overview dashboard” on page 132

Open the **Managed System Groups Overview** dashboard to see the consolidated situation event status for each managed system group in your environment.

Page layout and controls

Use the server dashboard tools for manipulating the page and accessing other dashboards.

Optimal window size for dashboards



The Infrastructure Management Dashboards for Servers are designed for a screen resolution of 1280 x 1024 pixels. For the best fit of dashboard elements, adjust your browser window to 1280 x 1024 pixels, or maximize the window if your monitor is at that resolution.

Resource navigation icons



Managed System Groups returns you to the home dashboard. (For details, see “Managed System Groups Overview dashboard” on page 132.)



Situation Events opens a table of situation events for all managed systems that report to the hub Tivoli Enterprise Monitoring Server. (For details, see “Situation Events dashboard” on page 135.) Notice that the alert indicator, such as  **Fatal** or  **Warning**, next to the flag icon is updated automatically as new events arrive to show the highest severity of the open events on your managed network.

Top of the page

Along the top of the page, you can see the path to this dashboard and click a hyperlink to return to a previous dashboard. For example, if you are displaying the details about a situation event, the path might be Home > Managed System Groups Overview > *MyGroup* > *MySystem:LZ* > *MySituationEvent*. To move back to the last dashboard, you click the *MySystem:LZ* hypertext link.


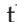
Do not use your browser's Back button to go to a previous dashboard. Dashboards are treated as a single portal page. Using the browser's Back button exits the application.

The **Actions** button opens a menu of options:

- Launching to the Tivoli Enterprise Portal (see “Launching to the Tivoli Enterprise Portal” on page 142).
- Copying the current URL for direct access by you or others that you send the link to, and for opening the page in a new tab or window (see “Copying the URL” on page 142).
- Setting the level of log tracing while you are using the dashboards (see “Server dashboard trace settings” in the *IBM Tivoli Monitoring Troubleshooting Guide*.)
- Displaying the server dashboards version information.

Managed systems groups and managed systems in a carousel or scorecard layout



The Managed System Groups Overview dashboard and the Overview tab of the individual managed system group dashboard both have a **Switch View** button to alternate between a carousel and scorecard layout of the managed system groups or managed systems. The scorecard view gives the same count of events by severity as the icon view, one row per managed system group. The scorecard also shows the number of managed systems in the group, how many managed systems are offline, and the name of the hub Tivoli Enterprise Monitoring Server that the managed systems in the group report to.

In the carousel view, if there are more groups or managed systems than can be seen in the view space, click the  and  to scroll through them.



In the scorecard view, click the count number for an event type to open the Situation Events dashboard.

Bar and line charts

Hover over a bar to display the value.

Click the  right or  left carousel button to cycle through the available charts.

Collapse and expand views

Click or drag the  bar on a view border to collapse or expand the view horizontally. Click or drag the  bar to collapse or expand the view vertically.

Exiting the dashboard application console

After you exit the console with the logout option, by closing the browser window, or by clicking the browser's Close button or Back button, it might take a moment for proper shutdown. Do not repeatedly click the Close or Back button or the browser might become unresponsive.

Related reference:

“Managed System Groups Overview dashboard” on page 132

Open the **Managed System Groups Overview** dashboard to see the consolidated situation event status for each managed system group in your environment.

“Managed system group dashboard” on page 133


Use the *managed system group* dashboard to get more information about the selected managed system group.

“Situation Events dashboard” on page 135

Use the Situation Events dashboard for an overview of all the open events in your managed network.

Table controls

Use the built-in table controls to reorganize the columns and rows and to hide those you are not interested in.

Filter Click inside the filter text box  and type the beginning of the value to filter the table by. As you type, the table rows that do not fit the criteria are filtered out and the **Total** is updated for the number of rows found.

Click the “x” in the filter box  or press the Backspace key to clear the filter.

Sort order

Click inside a column heading to sort by that column. Click the same column heading again to switch between ascending and descending sort order.

To add a secondary sort order, hold the Ctrl key down while you click the second column to sort by. You see a 1, 2, 3, and so on, next to the sort direction icon to indicate the sort value (primary, secondary, tertiary, and so on). For example, **1 ▲** is displayed for primary sort ascending and **2 ▼** is displayed for secondary sort descending.

Resize columns

Drag a column heading border to the right or left to adjust the column width.

Table pages

The table shows few rows per page, which ensures quick response time. To change the number of rows that are shown on the page, click one of the other values in the table status bar: 10, 25, 50, 100, or All. Click the **▲** to enter the page to display.

Click the left-arrow or right-arrow to move to the next page, or click a page number to go to that page.

Copying the URL

After you open another server dashboard, the URL in the browser address box does not change for the new location.

You can copy the unique URL for the dashboard that you are displaying and use it to open a new browser window with that dashboard or to access later by you or other users that you share the link with.

Procedure

1. Open the dashboard from which you want to copy the URL.
2. Click **Actions > Copy URL**.
3. Right-click the hypertext link (Link to the current page) and select one of the URL options to open the dashboard in a new browser tab or new window, or to copy the URL.

Launching to the Tivoli Enterprise Portal

You can launch to the Tivoli Enterprise Portal browser client to display more detailed metrics for a managed system and to perform such tasks as editing situations and acknowledging events.

For managed systems that are not supported in the server dashboards, you can launch to the browser client to open the default workspace for the managed system or attribute grouping.

Before you begin

Your user ID must be registered with the Tivoli Enterprise Portal Server. If single sign-on (SSO) has been established, you are not required to authenticate again when you launch to the Tivoli Enterprise Portal. For more information, see "User authentication through the portal server" in the *IBM Tivoli Monitoring Administrator's Guide*.

About this task

Complete these steps to open a new browser window and launch to the Tivoli Enterprise Portal to the default workspace for the current managed system type, individual managed system, or attribute grouping for that managed system. For example, if you launch the portal browser client from the Process tab of the Linux managed system dashboard, the Process workspace for that managed system is displayed.

Procedure

1. To supply the information that is needed to construct a URL for the workspace to open, select a managed system type or managed system in one of the server dashboards.
2. Click **Actions > Launch to TEP**.

Results

The Tivoli Enterprise Portal is started in your browser. Context information is passed from the dashboard to the Tivoli Enterprise Portal and the associated default workspace is displayed.

What to do next

If the Tivoli Enterprise Portal does not start, there might be a problem running the browser client for such reasons as the Java plug-in not being installed or configured correctly. You can review the browser client startup information in the *IBM Tivoli Monitoring Installation and Setup Guide* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/install/itm_install.htm) and browser client troubleshooting information in the *IBM Tivoli Monitoring Troubleshooting Guide* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/trouble/itm_troubleshoot.htm).

Setting a trace

IBM Infrastructure Management Dashboards for Servers web application has several levels of tracing that you can set while you work with dashboards. You can start a higher level of tracing exactly at the point in the user interface where you are having a problem, then return tracing to a lower level after capturing the necessary log data.


Adjust the trace settings to help your administrator or IBM Support to diagnose the cause of problems with the server dashboards.

For example, if a particular dashboard is behaving unexpectedly, you can raise the trace level before opening the dashboard to log the activity and then return trace logging to the normal level.

About this task

Take the following steps to set the trace level when you want to increase or reduce the amount of trace logging.

Procedure

1. If a server dashboard is not already open, select  **System Status and Health** > **Server Dashboards** and navigate to the dashboard where you want to change the trace level.
2. Click **Actions** > **Trace level** and select one of the following levels:
 - **Verbose** to have all activity logged. Verbose trace level includes Moderate, Light, and Minimal trace logging.
 - **Moderate** to have variable changes logged, such as what parameters were passed in and what calculations were made. Moderate trace level includes Light and Minimal trace logging.
 - **Light** to log error and variable activity. You might want to set the trace to this level if you have a problem such as no data being returned but the dashboard continues to function. Light trace level includes Minimal trace logging.
 - **Minimal** is the default setting and records only unrecoverable errors. You can set the trace level back to minimal after collecting a specific activity sequence. Even if a different trace level was set before logout, the trace is always reset to the lowest level the next time you log in.

Results

The trace is adjusted to the level chosen for this and all subsequent dashboards selected. To keep communications traffic to a minimum, the log messages are transferred in batches to the Dashboard Application Services Hub. A final transfer is made after you log out, whether manually or after a timeout period. (If the browser fails, no final logging is sent.)

The logs are saved on the server computer and named *userid.log.0* where *userid* is the ID used to log in to the Dashboard Application Services Hub and "0" is the first log. Three log files of 750 KB total are used to record trace data in a cyclical manner: After the *userid.log.0* reaches 250 KB, log entries are saved to *userid.log.1*; after *userid.log.1* reaches 250 KB, log entries go to *userid.log.2* until it reaches the maximum, at which time *userid.log.0* is cleared and new entries are saved there.

Chapter 9. Tivoli Common Reporting for the monitoring agent

This section 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.doc_211/ic-home.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.
- To execute reports from the command line using the **trcmd** command, the parameters **TCRDateRange** and **Forecast Period** are always required.

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 200 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.

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

Note that reports that are based on CCC Logs attributes, such as the Top Situations By Status report and the Situations History report, rely on raw data from the Status_History table in the Tivoli Data Warehouse. This table is populated by collecting historical data from the monitoring server for CCC Logs data after you select the Situation Status Log attribute group.

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

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.

Note: Within a chart, the label of an axis could be truncated depending on the chart size, chart scale, and specific metrics. This truncation is a restriction of the reporting tool.

Utilization Details for Single Resource report

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.

Characteristic	Description
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.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on, and is selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and is selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>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.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>Forecast Period If forecast is enabled, specifies the forecast period.</p> <p>Forecast Specifies whether forecast is enabled using a drop-down list. The list contains the following options:</p> <ul style="list-style-type: none"> • Use forecast • Do not use the forecast <p>Show Data Specifies if the chart data source should be displayed in a table or not.</p>

Characteristic	Description
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>CPU Utilization:</p> <p>Windows agent: NT_System_HV, NT_System_DV, NT_System_WV, NT_System_MV, NT_System_QV, NT_System_YV</p> <p>Linux agent: KLZ_CPU_HV, KLZ_CPU_DV, KLZ_CPU_WV, KLZ_CPU_MV, KLZ_CPU_QV, KLZ_CPU_YV</p> <p>UNIX agent: System_HV, System_DV, System_WV, System_MV, System_QV, System_YV</p> <p>Disk Utilization:</p> <p>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</p> <p>Linux agent: KLZ_Disk_HV, KLZ_Disk_DV, KLZ_Disk_WV, KLZ_Disk_MV, KLZ_Disk_QV, KLZ_Disk_YV</p> <p>UNIX agent: Disk_HV, Disk_DV, Disk_WV, Disk_MV, Disk_QV, Disk_YV</p> <p>Memory Utilization:</p> <p>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</p> <p>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</p> <p>UNIX agent: Unix_Memory_HV, Unix_Memory_DV, Unix_Memory_WV, Unix_Memory_MV, Unix_Memory_QV, Unix_Memory_YV</p> <p>Network Utilization:</p> <p>Windows agent: NT_Server_HV, NT_Server_DV, NT_Server_WV, NT_Server_MV, NT_Server_QV, NT_Server_YV</p> <p>Linux agent: KLZ_Network_HV, KLZ_Network_DV, KLZ_Network_WV, KLZ_Network_MV, KLZ_Network_QV, KLZ_Network_YV</p> <p>UNIX agent: Network_HV, NetworkDV, Network_WV, Network_MV, Network_QV, Network_YV</p> <p>Processes:</p> <p>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</p> <p>Linux agent: KLZ_Process_HV, KLZ_Process_DV, KLZ_Process_WV, KLZ_Process_MV, KLZ_Process_QV, KLZ_Process_YV</p> <p>UNIX agent: Process_HV, Process_DV, Process_WV, Process_MV, Process_QV, Process_YV</p>
Output	<p>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.</p>

Characteristic	Description
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	X
	Linux Disk	KLZ_Disk	X	X	X	X	X	X	X
	Linux VM Stats	KLZ_VM_Stats	X	X	X	X	X	X	X
	Linux Network	KLZ_Network	X	X	X	X	X	X	X
	Linux Process	KLZ_Process	X	X	X	X	X	X	X
UNIX	System	System	X	X	X	X	X	X	X
	Disk	Disk	X	X	X	X	X	X	X
	UNIX Memory	UNIX_Memory	X	X	X	X	X	X	X
	Network	Network	X	X	X	X	X	X	X
	Process	Process	X	X	X	X	X	X	X
Windows	System	NT_System	X	X	X	X	X	X	X
	Logical Disk	NT_Logical_Disk	X	X	X	X	X	X	X
	Memory	NT_Memory_64	X	X	X	X	X	X	X
	Server	NT_Server	X	X	X	X	X	X	X
	Process	NT_Process_64	X	X	X	X	X	X	X

Utilization Details for Multiple Resources report

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.

Characteristic	Description
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	<p>OS Type Determines the type of agent to work on and can be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>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.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days

Characteristic	Description
Parameters (Cont.)	<p>Forecast Period If forecast is enabled, specifies the forecast period.</p> <p>Forecast Specifies whether forecast is enabled using a drop-down list. The list contains the following options:</p> <ul style="list-style-type: none"> • Use forecast • Do not use the forecast
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>CPU Utilization:</p> <p>Windows agent: NT_System_HV, NT_System_DV, NT_System_WV, NT_System_MV, NT_System_QV, NT_System_YV</p> <p>Linux agent: KLZ_CPU_HV, KLZ_CPU_DV, KLZ_CPU_WV, KLZ_CPU_MV, KLZ_CPU_QV, KLZ_CPU_YV</p> <p>UNIX agent: System_HV, System_DV, System_WV, System_MV, System_QV, System_YV</p> <p>Disk Utilization:</p> <p>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</p> <p>Linux agent: KLZ_Disk_HV, KLZ_Disk_DV, KLZ_Disk_WV, KLZ_Disk_MV, KLZ_Disk_QV, KLZ_Disk_YV</p> <p>UNIX agent: Disk_HV, Disk_DV, Disk_WV, Disk_MV, Disk_QV, Disk_YV</p> <p>Memory Utilization:</p> <p>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</p> <p>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</p> <p>UNIX agent: Unix_Memory_HV, Unix_Memory_DV, Unix_Memory_WV, Unix_Memory_MV, Unix_Memory_QV, Unix_Memory_YV</p> <p>Network Utilization:</p> <p>Windows agent: NT_Server_HV, NT_Server_DV, NT_Server_WV, NT_Server_MV, NT_Server_QV, NT_Server_YV</p> <p>Linux agent: KLZ_Network_HV, KLZ_Network_DV, KLZ_Network_WV, KLZ_Network_MV, KLZ_Network_QV, KLZ_Network_YV</p> <p>UNIX agent: Network_HV, Network_DV, Network_WV, Network_MV, Network_QV, Network_YV</p>
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.

Characteristic	Description
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	X
	Linux Disk	KLZ_Disk	X	X	X	X	X	X	X
	Linux VM Stats	KLZ_VM_Stats	X	X	X	X	X	X	X
	Linux Network	KLZ_Network	X	X	X	X	X	X	X
UNIX	System	System	X	X	X	X	X	X	X
	Disk	Disk	X	X	X	X	X	X	X
	UNIX Memory	UNIX_Memory	X	X	X	X	X	X	X
	Network	Network	X	X	X	X	X	X	X
Windows	System	NT_System	X	X	X	X	X	X	X
	Logical Disk	NT_Logical_Disk	X	X	X	X	X	X	X
	Memory	NT_Memory_64	X	X	X	X	X	X	X
	Server	NT_Server	X	X	X	X	X	X	X

Utilization Comparison for Single Resource report

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.

Characteristic	Description
Purpose	This report helps to compare the CPU, disk, and memory utilization of a single server.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on, and is selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and is selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>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.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>Forecast Period If forecast is enabled, specifies the forecast period.</p> <p>Forecast Specifies whether forecast is enabled using a drop-down list. The list contains the following options:</p> <ul style="list-style-type: none"> • Use forecast • Do not use the forecast

Characteristic	Description
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>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</p> <p>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</p> <p>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</p>
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

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.

Characteristic	Description
Purpose	This report helps to compare the CPU, disk, and memory utilization for multiple servers.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on, and is selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and is selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>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.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>Forecast Period If forecast is enabled, specifies the forecast period.</p> <p>Forecast Specifies whether forecast is enabled using a drop-down list. The list contains the following options:</p> <ul style="list-style-type: none"> • Use forecast • Do not use the forecast

Characteristic	Description
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>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</p> <p>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</p> <p>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</p>
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

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.

Characteristic	Description
Purpose	Helps identify system performance of a system or server over a period of time. Shows daily patterns for utilization.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>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.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days

Characteristic	Description
Parameters (continued)	<p>Forecast Period If forecast is enabled, specifies the forecast period.</p> <p>Forecast Specifies whether forecast is enabled using a drop-down list. The list contains the following options:</p> <ul style="list-style-type: none"> • Use forecast • Do not use the forecast <p>Upper Limit for Good Status Specifies the upper limit for good status.</p> <p>Upper Limit for Fair Status Specifies the upper limit for fair status.</p> <p>Upper Limit for Warning Status Specifies the upper limit for warning status.</p> <p>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.</p>
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>CPU Utilization:</p> <p>Windows agent: NT_System_HV</p> <p>Linux agent: KLZ_CPU_HV</p> <p>UNIX agent: System_HV</p> <p>Disk Utilization:</p> <p>Windows agent: NT_Logical_Disk_HV</p> <p>Linux agent: KLZ_Disk_HV</p> <p>UNIX agent: Disk_HV</p> <p>Memory Utilization:</p> <p>Windows agent: NT_Memory_64_HV</p> <p>Linux agent: KLZ_VM_Stats_HV</p> <p>UNIX agent: Unix_Memory_HV</p>
Output	<p>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.</p>
Usage	<p>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.</p>
Drill through	<p>On row level to Utilization Details for Single Resource.</p>

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

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.

Characteristic	Description
Purpose	Helps identify which systems are performing poorly due to low physical memory causing excessive paging.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>Servers The server or system names for the selected OS Type is displayed in a drop-down list.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days

Characteristic	Description
Parameters (continued)	<p>Forecast Period If forecast is enabled, specifies the forecast period.</p> <p>Forecast Specifies whether forecast is enabled using a drop-down list. The list contains the following options:</p> <ul style="list-style-type: none"> • Use forecast • Do not use the forecast
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>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</p> <p>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</p> <p>UNIX agent: Unix_Memory_HV, Unix_Memory_DV, Unix_Memory_WV, Unix_Memory_MV, Unix_Memory_QV, Unix_Memory_YV</p>
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

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.

Characteristic	Description
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	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>Servers The server or system names for the selected OS Type is displayed in a drop-down list.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days

Characteristic	Description
Parameters (continued)	<p>Forecast Period If forecast is enabled, specifies the forecast period.</p> <p>Forecast Specifies whether forecast is enabled using a drop-down list. The list contains the following options:</p> <ul style="list-style-type: none"> • Use forecast • Do not use the forecast
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>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</p> <p>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</p> <p>UNIX agent: Unix_Memory_HV, Unix_Memory_DV, Unix_Memory_WV, Unix_Memory_MV, Unix_Memory_QV, Unix_Memory_YV</p>
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

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 eliminates over-crowding of the bars in the chart.

Characteristic	Description
Purpose	<p><i>CPU utilization:</i> 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.</p> <p><i>Disk utilization:</i> 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.</p> <p><i>Memory utilization:</i> 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.</p>

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>Number of systems The maximum number of systems to display.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>Resource A drop-down list that you can use to choose which type of resource to display:</p> <ul style="list-style-type: none"> • All • CPU • Disk • Memory

Characteristic	Description
Tables or views used	<p>CPU utilization</p> <ul style="list-style-type: none"> • 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 <p>Disk utilization</p> <ul style="list-style-type: none"> • 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 <p>Memory utilization</p> <ul style="list-style-type: none"> • 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	<p>A table is displayed with each row displaying a stacked bar representing one of the following for each system over the selected report period.</p> <ul style="list-style-type: none"> • 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) <p>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.</p>
Usage	<p>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.</p>
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				
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

OS Type	Attribute Group	Table	Summarization					
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

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.

Characteristic	Description
Purpose	Helps to analyze the top situations generating the selected event.
Parameters	<p>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:</p> <ul style="list-style-type: none"> • 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 <p>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:</p> <ul style="list-style-type: none"> • Acknowledged • Closed • Open • Reset • Stopped • Unknown <p>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.</p>
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.

Characteristic	Description
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

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.

Characteristic	Description
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

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 other 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.

Characteristic	Description
Purpose	This report helps to compare the CPU, disk and memory utilization of the machines with the same operating system in the Enterprise.
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>Date A date prompt where you can choose the date of the report.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>Attribute A drop-down list that you can use to choice what type of resource you would like to display:</p> <ul style="list-style-type: none"> • All (Default) • CPU • Disk • Memory <p>Upper Limit for Good Status Specifies the upper limit for good status.</p> <p>Upper Limit for Fair Status Specifies the upper limit for fair status.</p> <p>Upper Limit for Warning Status Specifies the upper limit for warning status.</p> <p>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.</p>
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>Windows agent: NT_System_HV, NT_Logical_Disk_HV, NT_Memory_64_HV</p> <p>Linux agent: KLZ_CPU_HV, KLZ_Disk_HV, KLZ_VM_Stats_HV</p> <p>UNIX agent: System_HV, Disk_HV, Unix_Memory_HV</p>
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.

Characteristic	Description
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

This report shows the overall availability and utilization of all Windows, Linux and UNIX monitoring agents.

Characteristic	Description
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.

Characteristic	Description
Parameters	<p>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:</p> <ul style="list-style-type: none"> • 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 <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>Windows agent: NT_System_DV, NT_Memory_64_DV, NT_Logical_Disk_DV</p> <p>Linux agent: KLZ_CPU_DV, KLZ_VM_Stats_DV, KLZ_Disk_DV, KLZ_System_Statistics_DV</p> <p>UNIX agent: System_DV, Disk_DV, Unix_Memory_DV</p>
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

This report displays availability of the top N systems based on System Up time over a period of time.

Characteristic	Description
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.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>Number of systems The maximum number of systems to display.</p> <p>Sort by A drop-down list that you can use to choose how the top N list is sorted:</p> <ul style="list-style-type: none"> • % Up Time • % Down Time
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>Windows agent: NT_System_DV</p> <p>Linux agent: KLZ_System_Statistics_DV</p> <p>UNIX agent: System_DV</p>
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).

Characteristic	Description
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

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.

Characteristic	Description
Purpose	This report helps to compare the top servers by CPU, disk, and memory utilization.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>Sorting Attribute A drop-down list that you can use to choice what type of resource you would like to display:</p> <ul style="list-style-type: none"> • CPU (Default) • Disk • Memory <p>Number of Systems The maximum number of servers to show in the report.</p> <p>Upper Limit for Good Status Specifies the upper limit for good status.</p> <p>Upper Limit for Fair Status Specifies the upper limit for fair status.</p> <p>Upper Limit for Warning Status Specifies the upper limit for warning status.</p> <p>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.</p>

Characteristic	Description
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>Windows agent: NT_System_HV, NT_Logical_Disk_HV, NT_Memory_64_HV</p> <p>Linux agent: KLZ_CPU_HV, KLZ_Disk_HV, KLZ_VM_Stats_HV</p> <p>UNIX agent: System_HV, Disk_HV, Unix_Memory_HV</p>
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 <i>Number of systems</i> 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)

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).

Characteristic	Description
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.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>Number of systems The maximum number of systems to display.</p> <p>Sort by A drop-down list that you can use to choose how the top N list is sorted:</p> <ul style="list-style-type: none"> • Mean Time To Recovery (Default) • Mean Time Before System Interruption
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>Windows agent: NT_System_DV</p> <p>Linux agent: KLZ_System_Statistics_DV</p> <p>UNIX agent: System_DV</p>
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

This report shows availability comparison between two or more servers.

Characteristic	Description
Purpose	Helps compare multiple systems based on availability.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>Windows agent: NT_System_DV</p> <p>Linux agent: KLZ_System_Statistics_DV</p> <p>UNIX agent: System_DV</p>
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

This report helps identify patterns of resource availability over a period of time.

Characteristic	Description
Purpose	Helps identify system performance of a system or server over a period of time. Shows daily patterns for availability or unavailability.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only <p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>Attribute A drop-down list that you can use to specify which pattern to display:</p> <ul style="list-style-type: none"> • % Up Time (Default) • % Down Time <p>Upper Limit for Good Status Specifies the upper limit for good status.</p> <p>Upper Limit for Fair Status Specifies the upper limit for fair status.</p> <p>Upper Limit for Warning Status Specifies the upper limit for warning status.</p> <p>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.</p>
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>Windows agent: NT_System_HV</p> <p>Linux agent: KLZ_System_Statistics_HV</p> <p>UNIX agent: System_DV</p>

Characteristic	Description
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

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 the 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.

Characteristic	Description
Purpose	Helps to compare different system CPU usage behaviors to identify excessive CPU utilization, unbalanced workloads or wrong configurations.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and is selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>Servers The server or system names for the selected OS Type are displayed in a drop-down list.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only

Characteristic	Description
Parameters (Continued)	<p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>Forecast Period If forecast is enabled, specifies the forecast period.</p> <p>Forecast Specifies whether forecast is enabled using a drop-down list. The list contains the following options:</p> <ul style="list-style-type: none"> • Use forecast • Do not use the forecast
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>Windows agent: NT_System_HV, NT_System_DV, NT_System_WV, NT_System_MV, NT_System_QV, NT_System_YV</p> <p>Linux agent: KLZ_CPU_HV, KLZ_CPU_DV, KLZ_CPU_WV, KLZ_CPU_MV, KLZ_CPU_QV, KLZ_CPU_YV</p> <p>UNIX agent: System_HV, System_DV, System_WV, System_MV, System_QV, System_YV</p>
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

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 the 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.

Characteristic	Description
Purpose	Helps identify which systems are experiencing excessive CPU usage.
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and is selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>Servers The server or system names for the selected OS Type are displayed in a drop-down list.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only

Characteristic	Description
Parameters (Continued)	<p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>Forecast Period If forecast is enabled, specifies the forecast period.</p> <p>Forecast Specifies whether forecast is enabled using a drop-down list. The list contains the following options:</p> <ul style="list-style-type: none"> • Use forecast • Do not use the forecast
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>Windows agent: NT_System_HV, NT_System_DV, NT_System_WV, NT_System_MV, NT_System_QV, NT_System_YV</p> <p>Linux agent: KLZ_CPU_HV, KLZ_CPU_DV, KLZ_CPU_WV, KLZ_CPU_MV, KLZ_CPU_QV, KLZ_CPU_YV</p> <p>UNIX agent: System_HV, System_DV, System_WV, System_MV, System_QV, System_YV</p>
Output	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

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 the 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.

Characteristic	Description
Purpose	Helps to analyze the disk utilization details of a specific machine.
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and is selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>Servers The server or system names for the selected OS Type are displayed in a drop-down list.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only

Characteristic	Description
Parameters (Continued)	<p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>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.</p> <p>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.</p> <p>Forecast Period If forecast is enabled, specifies the forecast period.</p> <p>Forecast Specifies whether forecast is enabled using a drop-down list. The list contains the following options:</p> <ul style="list-style-type: none"> • Use forecast • Do not use the forecast
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>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</p> <p>Linux agent: KLZ_Disk_HV, KLZ_Disk_DV, KLZ_Disk_WV, KLZ_Disk_MV, KLZ_Disk_QV, KLZ_Disk_YV</p> <p>UNIX agent: Disk_HV, Disk_DV, Disk_WV, Disk_MV, Disk_QV, Disk_YV</p>
Output	A line chart showing the average percent space usage plotted against time. A table showing finer disk utilization details.
Usage	<p>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.</p> <p>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.</p>
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 DISK	KLZ_DISK	X	X	X	X	X	X	X
Windows	Logical Disk	NT_Logical_Disk	X	X	X	X	X	X	X
UNIX	Disk	Disk	X	X	X	X	X	X	X

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. 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 the 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.

Characteristic	Description
Purpose	Helps to compare different file system usage behaviors to identify excessive file system utilization.

Characteristic	Description
Parameters	<p>OS Type Determines the type of agent to work on and should be selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Linux • UNIX • Windows <p>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:</p> <ul style="list-style-type: none"> • 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 <p>Summarization Type Determined by Summarization and Pruning and is selected from the drop-down list with the following items:</p> <ul style="list-style-type: none"> • Daily (Default) • Hourly • Weekly • Monthly • Quarterly • Yearly <p>Servers The server or system names for the selected OS Type are displayed in a drop-down list.</p> <p>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:</p> <ul style="list-style-type: none"> • All Shifts (Default) • Peak Hours Only • Off - Peak Hours Only

Characteristic	Description
Parameters (Continued)	<p>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:</p> <ul style="list-style-type: none"> • All Days (Default) • Work days • Vacation days <p>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.</p> <p>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.</p> <p>Forecast Period If forecast is enabled, specifies the forecast period.</p> <p>Forecast Specifies whether forecast is enabled using a drop-down list. The list contains the following options:</p> <ul style="list-style-type: none"> • Use forecast • Do not use the forecast
Tables or views used	<p>General: KSY_SUMMARIZATION_CONFIG_DV</p> <p>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</p> <p>Linux agent: KLZ_Disk_HV, KLZ_Disk_DV, KLZ_Disk_WV, KLZ_Disk_MV, KLZ_Disk_QV, KLZ_Disk_YV</p> <p>UNIX agent: Disk_HV, Disk_DV, Disk_WV, Disk_MV, Disk_QV, Disk_YV</p>
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	<p>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.</p> <p>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.</p>
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	X
Windows	Logical Disk	NT_Logical_Disk	X	X	X	X	X	X	X
UNIX	Disk	Disk	X	X	X	X	X	X	X

Situations History report

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.

Characteristic	Description
Purpose	You can use this report to analyze the history of the IBM Tivoli Monitoring situation events.

Characteristic	Description
Parameters	<p>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:</p> <ul style="list-style-type: none"> • 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 <p>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:</p> <ul style="list-style-type: none"> • Acknowledged • Closed • Open • Reset • Stopped • Unknown <p>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 (' ').</p> <p>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 (\).</p>
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.

Characteristic	Description
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 Table 3 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. Attribute 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. Attribute groups supported by the data model (continued)

Agent	Attribute groups	Tables
UNIX	Disk	Disk
	Network	Network
	Process	Process
	Unix Memory	Unix_Memory
	System	System
	AIX LPAR	AIX_LPAR
	AIX WPAR Information	AIX_WPAR_Information
	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
	Unix Print Queue	Unix_Print_Queue
	Unix Group	Unix_Group
	Unix Ping	Unix_Ping
	Unix All Users	Unix_All_Users
	Machine Information	Machine_Information
	Unix IP Address	Unix_IP_Address
	TCP Statistics	TCP_Statistics
	AIX AMS	AIX_AMS
	AIX Devices	AIX_Devices
	AIX WPAR CPU	AIX_WPAR_CPU
	AIX WPAR Physical Memory	AIX_WPAR_Physical_Memory
	AIX WPAR Network	AIX_WPAR_Network
	AIX WPAR File System	AIX_WPAR_FileSystem
	AIX Defined Users	AIX_Defined_Users
	AIX Physical Volumes	AIX_Physical_Volumes
AIX Volume Groups	AIX_Volume_Groups	
AIX Logical Volumes	AIX_Logical_Volumes	

Table 3. Attribute groups supported by the data model (continued)

Agent	Attribute groups	Tables
Windows	Logical Disk	NT_Logical_Disk
	Memory	NT_Memory_64
	Network Interface	NT_Network_Interface
	Process	NT_Process_64
	Server	NT_Server
	System	NT_Process_64
	ICMP Statistics	ICMP_Statistics
	IP Statistics	IP_Statistics
	Cache	NT_Cache
	Device Dependencies	NT_Device_Dependencies
	Devices	NT_Devices
	Event Log	NT_Event_Log
	Monitored Logs Report	NT_Monitored_Logs_Report
	Network Port	NT_Network_Port
	Objects	NT_Objects
	Paging File	NT_Paging_File
	Physical Disk	NT_Physical_Disk
	Printer	NT_Printer
	Processor	NT_Processor
	Processor Summary	NT_Processor_Summary
	Redirector	NT_Redirector
	Server Work Queues	NT_Server_Work_Queues_64
	Service Dependencies	NT_Service_Dependencies
	Services	NT_Services
	Thread	NT_Thread
	Print Queue	Print_Queue
	Process IO	Process_IO
TCP Statistics	TCP_Statistics	
UDP Statistics	UDP_Statistics	

Table 3. Attribute groups supported by the data model (continued)

Agent	Attribute groups	Tables
IBM i	OS400 System Status	OS400_System_Status
	OS400 Acct Jrn	OS400_Acct_Jrn
	OS400 Alert	OS400_Alert
	OS400 APPN Topology	OS400_APPN_Topology
	OS400 Comm Async	OS400_Comm_Async
	OS400 Comm Bisync	OS400_Comm_Bisync
	OS400 Controller	OS400_Controller
	OS400 DB Member	OS400_DB_Member
	OS400 Device	OS400_Device
	OS400 Disk Unit	OS400_Disk_Unit
	OS400 Comm Ethernet	OS400_Comm_Ethernet
	OS400 Job Queue	OS400_Job_Queue
	OS400 Line	OS400_Line
	OS400 Message	OS400_Message
	OS400 Network	OS400_Network
	OS400 Object	OS400_Object
	OS400 I/O Processor	OS400_I/O_Processor
	OS400 Job	OS400_Job
	OS400 Storage Pool	OS400_Storage_Pool
	OS400 Subsystem	OS400_Subsystem
	OS400 Comm SDLC	OS400_Comm_SDLC
	OS400 Security Jrn AuthFail	OS400_Security_Jrn_AuthFail
	OS400 Security Jrn AuditJrn	OS400_Security_Jrn_AuditJrn
OS400 Security Jrn ChgAuth	OS400_Security_Jrn_ChgAuth	
OS400 Security Jrn ChgUseProf	OS400_Security_Jrn_ChgUseProf	

Table 3. Attribute groups supported by the data model (continued)

Agent	Attribute groups	Tables
IBM i (continued)	OS400 Security Jrn JobDesc	OS400_Security_Jrn_JobDesc
	OS400 Security Jrn Network	OS400_Security_Jrn_Network
	OS400 Security Jrn ChgOwner	OS400_Security_Jrn_ChgOwner
	OS400 Security Jrn ProgAdopt	OS400_Security_Jrn_ProgAdopt
	OS400 Security Jrn ProfSwap	OS400_Security_Jrn_ProfSwap
	OS400 Security Jrn Password	OS400_Security_Jrn_Password
	OS400 Security Jrn RestoreJob	OS400_Security_Jrn_RestoreJob
	OS400 Security Jrn RestoreProg	OS400_Security_Jrn_RestoreProg
	OS400 Security Jrn SYSVAL	OS400_Security_Jrn_SYSVAL
	OS400 System Values Acct	OS400_System_Values_Acct
	OS400 System Values	OS400_System_Values
	OS400 System Values Device	OS400_System_Values_Device
	OS400 System Values IPL	OS400_System_Values_IPL
	OS400 System Values Prob	OS400_System_Values_Prob
	OS400 System Values Perf	OS400_System_Values_Perf
	OS400 System Values User	OS400_System_Values_User
	OS400 System Values Ring	OS400_System_Values_Ring
	OS400 Comm X25	OS400_Comm_X25
	i5OS Auxiliary Storage Pool	i5OS_Auxiliary_Storage_Pool
	i5OS TCPIP Logical Interface	i5OS_TCPIP_Logical_Interface
	i5OS TCPIP Service	i5OS_TCPIP_Service
	i5OS Network Interface	i5OS_Network_Interface
	i5OS Network Server	i5OS_Network_Server
	i5OS System Statistics	i5OS_System_Statistics
	i5OS Disk	i5OS_Disk
	i5OS Output Queue	i5OS_Output_Queue

Table 3. Attribute groups supported by the data model (continued)

Agent	Attribute groups	Tables
IBM i (continued)	i5OS History Log	i5OS_History_Log
	i5OS Integrated File System Object	i5OS_Integrated_File_System_Object
	i5OS Job Log	i5OS_Job_Log
	i5OS Net Server	i5OS_Net_Server
	i5OS Management Central	i5OS_Management_Central
	i5OS Distribution Queue	i5OS_Distribution_Queue
	i5OS Miscellaneous	i5OS_Miscellaneous
	i5OS Inactive Job	i5OS_Inactive_Job
	i5OS User and Group	i5OS_User_and_Group
	i5OS System Value Sys Ctl 1	i5OS_System_Value_Sys_Ctl_1
	i5OS System Value Sys Ctl 2	i5OS_System_Value_Sys_Ctl_2
	i5OS System Value Allocation	i5OS_System_Value_Allocation
	i5OS System Value Date Time	i5OS_System_Value_Date_Time
	i5OS System Value Editing	i5OS_System_Value_Editing
	i5OS System Value Security	i5OS_System_Value_Security
	i5OS System Value Other	i5OS_System_Value_Other
	i5OS TCPIP Route	i5OS_TCPIP_Route
	i5OS TCPIP Host	i5OS_TCPIP_Host
	i5OS Cluster Node	i5OS_Cluster_Node
	i5OS Cluster Resource Group	i5OS_Cluster_Resource_Group
	i5OS Cluster Monitored Resources	i5OS_Cluster_Monitored_Resources
	i5OS Licensed Program Product	i5OS_Licensed_Program_Product
	i5OS Program Temporary Fix	i5OS_Program_Temporary_Fix
i5OS Group Program Temporary Fix	i5OS_Group_Program_Temporary_Fix	
i5OS Group PTF Details	i5OS_Group_PTF_Details	
i5OS IOA Cache Battery	i5OS_IOA_Cache_Battery	
CCC Logs	Situation Status Log	TSITSTSH

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 10. Troubleshooting

This chapter explains how to troubleshoot the IBM Tivoli Monitoring: Linux 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 5.

This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information. Also see “Support information” on page 311 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 208 for lists of all trace log files and their locations. See the <i>IBM Tivoli Monitoring User’s Guide</i> for general information about the IBM Tivoli Monitoring environment.
Linux information	<ul style="list-style-type: none">• 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: <ul style="list-style-type: none">• IBM Tivoli Monitoring. Also provide the patch level, if available.• IBM Tivoli Monitoring: Linux OS Agent
Screen captures	Screen captures of incorrect output, if any.
Core dump files	If the system stops on UNIX or Linux systems, collect core dump file from <i>install_dir/bin</i> directory, where <i>install_dir</i> is the directory path where you installed the monitoring agent.

Built-in troubleshooting features

The primary troubleshooting feature in the IBM Tivoli Monitoring: Linux OS Agent is logging. *Logging* refers to the text messages and trace data generated by the IBM Tivoli Monitoring: Linux 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 208 for more information.

Problem classification

The following types of problems might occur with the IBM Tivoli Monitoring: Linux 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 209
- “Examples: using trace logs” on page 210
- “Setting RAS trace parameters” on page 211

Note: The documentation refers to the RAS facility in IBM Tivoli Monitoring as “RAS1”.

IBM Software Support uses the information captured by trace logging to trace a problem to its source or to determine why an error occurred. 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.

Log file management is described in the following table:

Table 5. Log file management on UNIX compared to log file management on Windows

Location of logs	Description
<ul style="list-style-type: none">• On a Windows monitoring server• On a Windows computer where the monitoring agent is running• On a UNIX or Linux computer where the monitoring agent is running	<p>On Windows, the log file is overwritten each time the component starts. There is no automated method to archive previous RAS1 log files.</p> <p>Note: To prevent the log files from consuming too much disk space, you can stop and start the component. This action automatically creates a new log file. Save a backup of log files if your company policy requires archiving of log files.</p>

Table 5. Log file management on UNIX compared to log file management on Windows (continued)

Location of logs	Description
<ul style="list-style-type: none"> On a UNIX or Linux monitoring server On a UNIX or Linux computer where the monitoring agent is running 	<p>On UNIX or Linux systems, because of the use of the <code>&Timestamp;</code> variable in the log file names, multiple RAS1 logs are normally stored the logs subdirectory. The file name for a trace log is a copy of a related file that includes the process ID of the agent. The two files have the same timestamp as in these examples from a computer with a host name f50pa2b. The 1112097194 part of the name is the timestamp:</p> <pre>f50pa2b_lz_1112097194.log f50pa2b_lz_1112097194.pid60420</pre> <p>where <i>lz</i> is the unique, two-character code for Monitoring Agent for Linux OS.</p>

Note: When you communicate with IBM Software Support, you must capture and send the RAS1 log that matches any problem occurrence that you report. Table 6 can help you identify files that are relevant to your troubleshooting efforts.

Principal trace log files

Table 6 contains locations, file names, and descriptions of trace logs that can help determine the source of problems with agents.

Table 6. Trace log files for troubleshooting agents

System where log is located	File name and path	Description
On the computer that hosts the monitoring agent	The <i>hostname_lz_instance.log</i> file is located in the <i>install_dir/logs</i> path.	Traces activity of the monitoring agent.
	The *.LG0 file is located in the following subdirectory of the <i>install_dir</i> path: <i>/logs</i> .	Shows whether agent was able to connect to the monitoring server. Shows which situations are started and stopped, and shows other events while the agent is running. 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 .LG1 to learn the following details regarding the <i>previous</i> monitoring session: <ul style="list-style-type: none"> Status of connectivity with the monitoring server. Situations that were running. The success or failure status of Take Action commands.
	The <i>take_action_name.log</i> file (where <i>take_action_name</i> is the name of the Take Action command) is located in the <i>/logs</i> subdirectory of the <i>install_dir</i> path.	Traces activity each time a Take Action command runs. For example, when a hypothetical start_command Take Action command runs, IBM Tivoli Monitoring would generate a <i>start_command.log</i> file.

Table 6. Trace log files for troubleshooting agents (continued)

System where log is located	File name and path	Description
On the Tivoli Enterprise Monitoring Server	The <code>install_dir/logs</code> file in the <code>install_dir/logs</code> 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 <code>Warehouse_Configuration.log</code> file is located in the following path on Windows: <code>install_dir\InstallITM</code> .	Provides details about the configuration of data warehousing for historical reporting.
	The name of the RAS log file is as follows: <ul style="list-style-type: none"> • On Windows: <code>install_dir\logs\hostname_ms_timestamp.log</code> • On UNIX or Linux: <code>hostname_ms_timestamp.log</code> and <code>hostname_ms_timestamp.pidnnnn</code> in the <code>install_dir/logs</code> path, where <code>nnnnn</code> is the process ID number. 	Traces activity on the monitoring server.
On the Tivoli Enterprise Portal Server	The name of the RAS log file is as follows: <ul style="list-style-type: none"> • On Windows: <code>install_dir\logs\hostname_cq_timestamp.log</code> • On UNIX or Linux: <code>hostname_cq_timestamp.log</code> and <code>hostname_cq_timestamp.pidnnnn</code> in the <code>install_dir/logs</code> path, where <code>nnnnn</code> is the process ID number. 	Traces activity on the portal server.
	The <code>TEPS_ODBC.log</code> file is located in the following path on Windows: <code>install_dir\InstallITM</code> .	When you enable historical reporting, this log file traces the status of the warehouse proxy agent.
<p>Definitions of variables:</p> <p><code>timestamp</code> is timestamp whose format includes year (y), month (m), day (d), hour (h), and minute (m), as follows: yyymmdd hhmm</p> <p><code>install_dir</code> represents the directory path where you installed the IBM Tivoli Monitoring component.</p> <p><code>install_dir</code> can represent a path on the computer that host the monitoring system, the monitoring agent, or the portal.</p> <p><code>instance</code> refers to the name of the database instance that you are monitoring.</p> <p><code>hostname</code> refers to the name of the computer on which the IBM Tivoli Monitoring component runs.</p>		

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 to learn some basic facts about your IBM Tivoli Monitoring environment.

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}kdc10c1.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:LZ is ON-LINE.
```

```
(42C3079B.0000-6A4:kpxreqhb.cpp,644,"HeartbeatInserter") Remote node SERVER5B:KLZ is OFF-LINE.
```

Key points regarding the preceding excerpt:

- The monitoring server appends the **LZ** product code to the server name to form a unique name (SERVER5B:LZ) for this instance of Monitoring Agent for Linux 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 Linux OS uses RAS1 tracing and generates the logs described in Table 6 on page 209. The default RAS1 trace level is ERROR.

Before you begin

When you are troubleshooting, follow these guidelines to ensure that you capture and analyze the correct log files: Because of the use of the `&Timestamp;` variable in the log file names on UNIX or Linux systems, there are typically multiple RAS1 logs in the logs subdirectory. When you forward log files to IBM Software Support, you must send the RAS1 log that matches the problem occurrence that the log files are reporting.

After you finish

On UNIX or Linux, periodically prune the trace logs in the logs subdirectory so that there is available disk space for new logging.

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/lz.ini* file. The basic format for setting tracing options is as follows:

```
KBB_RAS1=ERROR (UNIT:k1z options)
```

Use one of the following methods to modify trace options:

- **Manually edit the configuration file to set trace logging**
 1. Open the trace options file: *install_dir/config/lz.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:k1z ALL) (UNIT:kra ALL)'
```
 3. Restart the monitoring agent so that your changes take effect.

Problems and workarounds

This publication provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

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 5.

The following sections provide symptoms and workarounds for problems that might occur with Monitoring Agent for Linux OS:

- “Installation and configuration troubleshooting”
- “Agent troubleshooting” on page 217
- “Tivoli Enterprise Portal troubleshooting” on page 219
- “Troubleshooting for remote deployment” on page 220
- “Tivoli Common Reporting troubleshooting” on page 220
- “Situation troubleshooting” on page 221

Installation and configuration troubleshooting

This section provides tables that show solutions for installation, configuration, and uninstallation problems.

Table 7. Problems and solutions for installation and configuration

Problem	Solution
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.

Table 7. Problems and solutions for installation and configuration (continued)

Problem	Solution
install.sh fails with a JVMJ9VM011W error.	The SELINUX parameter in the /etc/sysconfig/selinux file must be set to "disable". Then, reboot the system.
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.
Installation of Monitoring Agent for Linux OS on the Linux S390 R2.6 64-bit operating system fails with a message similar to the following: LINUX MONITORING AGENT V610R <i>nnn</i> unable to install agent, where <i>nnn</i> is the release number.	Solve this problem as follows: <ol style="list-style-type: none"> 1. Run the following command before running any installation or configuration command for the agent: <pre>export JAVA_COMPILER=NONE</pre> 2. Install the following two RPM (Red Hat Package Manager) files: <ul style="list-style-type: none"> • compat-libstdc++-295-2.....s390x.rpm • compat-libstdc++-33-3.....s390x.rpm It requires the two s390x.rpm files, in addition to the s390.rpm files. <p>You can obtain the required RPM files from the CD for Red Hat As 4.0 s390x.</p>
During a command-line installation, you choose to install a component that is already installed, and you see the following warning: <pre>WARNING - you are about to install the SAME version of "component"</pre> <p>where <i>component</i> is the name of the component that you are attempting to install. Note: This problem affects UNIX command-line installations. If you monitor only Windows environments, you would see this problem if you choose to install a product component (for example, a monitoring server) on UNIX.</p>	You must exit and restart the installation process. You cannot return to the list where you selected components to install. When you run the installer again, do not attempt to install any component that is already installed.
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: <pre>install.sh warning: unarchive of "/cdrom/unix/cienv1.tar" may have failed</pre>	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 <i>all</i> 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.
Cannot locate the KDCB0_HOSTNAME setting.	Go to <i>install_dir/config</i> and edit the corresponding .ini file. Set the KDCB0_HOSTNAME parameter followed by the IP address. If you use multiple network interface cards (NICs), give the Primary IP address of the network interface.
The Monitoring Agent for Linux OS repeatedly restarts.	You can collect data to analyze this problem as follows: <ol style="list-style-type: none"> 1. Access the <i>install_dir/config/lz.ini</i> file, which is described in "Setting RAS trace parameters" on page 211. 2. Add the following line: KBB_SIG1=trace -dumpoff

Table 7. Problems and solutions for installation and configuration (continued)

Problem	Solution
<p>Agents in the monitoring environment use different communication protocols. For example, some agents have security enabled and others do not.</p>	<p>Configure both the monitoring server and the Warehouse proxy server to accept multiple protocols, as described in the <i>IBM Tivoli Monitoring Installation and Setup Guide</i>.</p>
<p>Creating a firewall partition file: The partition file enables an agent to connect to the monitoring server through a firewall.</p>	<p>How it works: When the agents start, they search KDCPARTITION.TXT for the following matches:</p> <ul style="list-style-type: none"> • An entry that matches the partition name OUTSIDE. • An entry that also includes a valid external address. <p>For more information, see the <i>IBM Tivoli Monitoring Installation and Setup Guide</i>.</p>
<p>You see the following error: Hub not registered with location broker. Error-code 1195.</p>	<p>Confirm that the password within the Tivoli Enterprise Monitoring Server is correct.</p>
<p>The Monitoring Agent for Linux OS is started and running but not displaying data in the Tivoli Enterprise Portal.</p>	<p>Perform the following steps:</p> <ol style="list-style-type: none"> 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 in <i>IBM Tivoli Monitoring Installation and Setup Guide</i> for more information on adding application support. 4. Check the log files to see whether there are connection problems. 5. If there are no connection problems, check whether the agent has terminated. 6. If the agent is not terminated, confirm that you have added application support for the Monitoring Agent for Linux OS in the Tivoli Enterprise Monitoring Server as follows: <ul style="list-style-type: none"> • Verify that the following entries are available in the <i>install_dir\candle_installation.log</i> file: <i>install_dir\Install\IBM Tivoli Monitoring timestamp.log</i> ... Browser Client support for ITM Agent for Linux ... Desktop Client support for ITM Agent for Linux <p>If the candle_installation.log file does not have the above entries for Monitoring Agent for Linux OS, add application support for this monitoring agent. See in <i>IBM Tivoli Monitoring Installation and Setup Guide</i> for more information on adding application support.</p> <ul style="list-style-type: none"> • Verify that the following files are available in the directory: <pre>install_dir\ATTRLIB\klz.atrinstall_dir \CNPS\CMSATR\klz.atrinstall_dir \SQLLIB\klz.sqlinstall_dir \CNPS\SQLLIB\klz.sql</pre>

Table 7. Problems and solutions for installation and configuration (continued)

Problem	Solution
You successfully migrate an OMEGAMON monitoring agent to IBM Tivoli Monitoring, Version 6.2.0. 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.	<p>Install the agent's application support files on the Tivoli Enterprise Monitoring Server, by using the following steps:</p> <ol style="list-style-type: none"> 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 in IBM Tivoli Monitoring Installation and Setup Guide for more information on adding application support. <p>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.</p>
You receive the following error: /data/itm/li6263/lz/bin/klzagent: error while loading shared libraries: libstdc++.so.5: cannot open shared object file: No such file or directory	Ensure that the libstdc++.so.5 library is installed.

Table 8. 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.	<p>When you want to remove a managed system from the navigation tree, complete the following steps:</p> <ol style="list-style-type: none"> 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.

Agent upgrade and restart using non-root

The monitoring agent can run by 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 by using the **Run As** option on the GUI or by using the **_UNIX_STARTUP_Username** option on the **tacmd addSystem** command line. If the agent is running by 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 by 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 by 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 by 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.itmNUM*
- HP-UX: */sbin/init.d/ITMAgentsNUM*
- Linux: */etc/init.d/ITMAgentsNUM*
- Solaris: */etc/init.d/ITMAgentsNUM*

3. Add entries for your operating system by 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, lz for the Monitoring Agent for Linux OS

Examples:

• For AIX, add entries with the following format:

```
su - USER -c " /opt/IBM/ITM/bin/itmcmd agent
-o INSTANCE start lz"
```

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 lz >/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.

Unique names for monitoring components

IBM Tivoli Monitoring might not be able to generate a unique name for monitoring components due to the truncation of names that the product automatically generates.

IBM Tivoli Monitoring automatically creates a name for each monitoring component by concatenating the host name and product code separated by colons (*hostname:LZ*).

Note: When you monitor a multinode system, such as a database, IBM Tivoli Monitoring adds a subsystem name to the concatenated name, typically a database instance name.

The length of the name that IBM Tivoli Monitoring generates is limited to 32 characters. Truncation can result in multiple components having the same 32-character name. If this problem happens, shorten the *hostname* portion of the name as follows:

1. Open the configuration file for the monitoring agent, which is located in the following path: *install_dir/config/lz.ini*.
2. Find the line that 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 *LZ*, 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 cannot find the **CTIRA_HOSTNAME** environment variable, you must add it to the configuration file of the monitoring agent:

- **On UNIX and Linux:** Add the variable to the *config/product_code.ini* file.

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 9. Agent problems and solutions

Problem	Solution
<p>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.</p>	<p>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.</p> <p>"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).</p> <p>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 by using the 'PORT:' keyword on the KDC_FAMILIES / KDE_TRANSPORT environment variable and defaults to '1918'.)</p> <p>The physical port allocation method is defined as $(BASE_PORT + 4096 * N)$ where $N=0$ for a Tivoli Enterprise Monitoring Server process and $N=\{1, 2, \dots, 15\}$ for a non-Tivoli Enterprise Monitoring Server. Two architectural limits result as a consequence of the physical port allocation method:</p> <ul style="list-style-type: none"> • 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 than 15 IP.PIPE processes can be active on a single system image. <p>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.</p> <p>No more than 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 than 15 agents per system image.</p> <p>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 by 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.)</p>
<p>The Monitoring Agent for Linux OS running on a Linux system does not communicate with the Tivoli Enterprise Monitoring Server running on a Z/OS system.</p>	<p>The procedure for seeding the Tivoli Enterprise Monitoring Server running on a Z/OS system for an instance of the Monitoring Agent for Linux OS running on a Linux system can be found in <i>Configuring Tivoli Enterprise Monitoring Server on z/OS®</i>.</p>

Table 9. Agent problems and solutions (continued)

Problem	Solution
<p>The agent's process, klzagent uses a large amount of system resources.</p>	<p>In most cases, the problem occurs during the backup. Any one of the following scenarios can cause this problem.</p> <p>The agent is running during the backup After backing up, the agent is started during system startup.</p> <p>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.</p> <p>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 Linux 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.</p> <p>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.</p> <p>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 log file has the following information: Unable to find running Tivoli Enterprise Monitoring Server on CMSLIST</p>
<p>Attributes do not allow non-ASCII input in the situation editor.</p>	<p>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.</p>
<p>In the User workspace, data does not show up in the User Login Information (table view).</p>	<p>This problem arises when you install the agent on a 64-bit zLinux operating system, but run the agent in 32-bit mode. The workspace is unable to access user login data. Run the agent in 64-bit mode.</p>
<p>Query produces no historical data</p>	<p>If you use wildcards within a query, the value does not act as a wildcard against historical data. It acts simply as a value to compare against the values in the historical table. However, the value does act as a wildcard against the realtime data.</p> <p>For example, if you use <code>.*(LongDirName/sleep).*</code> in the historical collection configuration and use <code>.*(LongDirName/sleep).*</code> in a query as well, you will see real time data and historical data. But, if you use <code>.*(LongDirName/sleep).*</code> in the historical collection configuration, and then use <code>.*(LongDirName/sle).*</code> in the query, then you will see real time data only and no historical data. The value does not act as a wildcard against historical data.</p>

Tivoli Enterprise Portal troubleshooting

The following table lists problems that might occur with the Tivoli Enterprise Portal. This section provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Table 10. Tivoli Enterprise Portal problems and solutions

Problem	Solution
Historical data collection is unavailable because of incorrect queries in the Tivoli Enterprise Portal.	<p>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.</p> <p>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).</p> <p>To ensure support of historical data collection, do not use the Sort By, Group By, or First/Last functions in your queries.</p> <p>See the <i>IBM Tivoli Monitoring Administrator's Guide</i> or the Tivoli Enterprise Portal online Help for information on the Historical Data Collection function.</p>
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.
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.

Troubleshooting for remote deployment

Table 11 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 11. 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.

Tivoli Common Reporting troubleshooting

Table 12 on page 221 lists general problems that might occur with Tivoli Common Reporting. This chapter provides agent-specific troubleshooting information. See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Table 12. Tivoli Common Reporting problems and solutions

Problem	Solution
<p>Reports that are based on CCC Logs attributes, such as the Top Situations By Status report and the Situations History report, rely on raw data from the Status_History table in the Tivoli Data Warehouse. To ensure accurate and complete reports, you might need to confirm historical collection of CCC Logs data.</p>	<p>The Status_History table is populated by collecting historical data from the monitoring server for CCC Logs data after you select the Situation Status Log attribute group. Historical collection of CCC Logs data can be confirmed by ensuring that the UADVISOR situation for O4SRV_TSITSTSH appears in the TEMS messages file.</p> <p>On Windows, the kdsmain.msg file in the: %CANDLE_HOME%\CMS\ path.</p> <p>On UNIX and Linux, the <hostname>_ms_<Timestamp>.log file in the <InstallDirectory>/logs/ path.</p> <p>Example entries in the TEMS messages log showing the O4SRV_TSITSTSH attribute gathered in a UADVISOR situation include:</p> <ul style="list-style-type: none"> • KO46256 Situation definition UADVISOR_O4SRV_TSITSTSH created by *ENTERPRISE. • KO41047 Situation UADVISOR_O4SRV_TSITSTSH distribution HUB_TEMS1 added. • KO41046 Monitoring for enterprise situation UADVISOR_O4SRV_TSITSTSH started. <p>After historical data is collected and exported and following the next warehousing interval of the TDW database, the Status_History table will be populated with data required to run the TCR reports using CCC Logs.</p>

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.

General situation problems

Table 13 lists problems that might occur with specific situations.

Table 13. Specific situation problems and solutions

Problem	Solution
<p>You want to change the appearance of situations when they are displayed in a Workspace view.</p>	<ol style="list-style-type: none"> 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.

Table 13. Specific situation problems and solutions (continued)

Problem	Solution
<p>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.</p>	<p>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.</p> <p>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.</p>
<p>Monitoring activity requires too much disk space.</p>	<p>Check the RAS trace logging settings that are described in “Setting RAS trace parameters” on page 211. For example, trace logs grow rapidly when you apply the ALL logging option.</p>
<p>A formula that uses mathematical operators appears 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 <code>#'Linux_VM_Stats.Total_Memory' / 10</code></p>	<p>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.</p>
<p>If you are running a Version 350 Monitoring Agent for Linux 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.</p>	<p>To enable Unicode and other features, upgrade the monitoring agent to IBM Tivoli Monitoring, Version 6.1.0.</p>
<p>IBM Tivoli Monitoring is configured to provide data to the optional product IBM Tivoli Enterprise Console. However, a situation displays the severity UNKNOWN in IBM Tivoli Enterprise Console.</p>	<p>For a situation to have the correct severity in TEC for those situations which are not mapped, you need to ensure that one of the following is true:</p> <ul style="list-style-type: none"> • Specify the severity in the SITINFO column of the O4SRV.TSITDESC table. For example use the values 'SEV=Critical' and 'SEV=Warning' for the SITINFO column in your kxx.sql file, which adds application support to the monitoring product. —OR— • Have the name of the situation ends with '_Warn' or '_Warning' for WARNING severity and '_Cri' or '_Critical' for Critical severity
<p>You see the 'Unable to get attribute name' error in the Tivoli Enterprise Monitoring Server log after creating a situation.</p>	<p>Install the agent's application support files on the Tivoli Enterprise Monitoring Server, by using the following steps:</p> <ol style="list-style-type: none"> 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 in IBM Tivoli Monitoring Installation and Setup Guide for more information on adding application support.
<p>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.</p>	<p>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.</p>

Table 13. Specific situation problems and solutions (continued)

Problem	Solution
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: <ol style="list-style-type: none"> 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 <i>CANDLEHOME/CMS/TECLIB/itm5migr</i> directory.
You are receiving Tivoli Business Systems Management events that cannot be associated due to <i>application_oid</i> and <i>application_class</i> 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 <i>agent_name_forward_tbsm_event_cb.sh</i> script on the Tivoli Enterprise Console server with the version of this file from the Tivoli Enterprise Monitoring Server in the <i>CANDLEHOME/CMS/TECLIB/itm5migr</i> directory.
Situations created using the File Pattern attribute group are always TRUE if the Match Count attribute is not used.	When creating situations using the File Pattern Attribute Group like the following one: <pre>[*IF *VALUE Linux_File_Pattern.File_Name *EQ '/path/filename' *AND *VALUE Linux_File_Pattern.Match_Pattern *EQ 'pattern']</pre> <p>the situations are always TRUE.</p> <p>Redefine the situations in the following way:</p> <pre>[*IF *VALUE Linux_File_Pattern.File_Name *EQ '/path/filename' *AND *VALUE Linux_File_Pattern.Match_Pattern *EQ 'pattern' *AND *VALUE Linux_File_Pattern.Match_Count *GT 0]</pre> <p>The Match_Count attribute must always be used in order to monitor for a match pattern.</p>

Problems with configuration of situations

Table 14 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 14. Problems with configuring situations that you solve in the Situation Editor

Problem	Solution
<p>Note: To get started with the solutions in this section, perform these steps:</p> <ol style="list-style-type: none"> 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 Linux OS has been added to the monitoring server. If not, add application support to the server, as described in the <i>IBM Tivoli Monitoring Installation and Setup Guide</i> .
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.

Table 14. Problems with configuring situations that you solve in the Situation Editor (continued)

Problem	Solution
The situation did not activate at startup.	Manually recycle the situation as follows: <ol style="list-style-type: none"> 1. Right-click the situation and choose Stop Situation. 2. Right-click the situation and choose Start Situation. <p>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.</p>
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: <ol style="list-style-type: none"> 1. Click the <i>fx</i> icon in the upper-right corner of the Formula area. The Show formula window is displayed. <ol style="list-style-type: none"> a. Confirm the following details in the Formula area at the top of the window: <ul style="list-style-type: none"> • 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. <p>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.</p>

Table 15. 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. <p>Note: The situation does not need to be displayed in the workspace. It is sufficient that the situation be associated with any workspace.</p>
You do not have access to a situation.	Note: You must have administrator privileges to perform these steps. <ol style="list-style-type: none"> 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.

Table 15. Problems with configuration of situations that you solve in the Workspace area (continued)

Problem	Solution
A managed system seems to be offline.	<ol style="list-style-type: none"> 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 16. 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.	Check the system status and check the appropriate IBM Tivoli Monitoring logs.
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.

Appendix A. Upgrading for warehouse summarization

The Monitoring Agent for Linux 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 only available from IBM Tivoli Monitoring v6.1.0 to v6.2.1, and is not available for upgrading from IBM Tivoli Monitoring v6.2 to v6.2.1.

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, `lnxallusr`.

- 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, `lnxallusr_H`, `lnxallusr_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 17 contains a list of the time periods and the suffixes for the summary tables and views.

Table 17. 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

Table 17. Time periods and suffixes for summary tables and views (continued)

Data collection time period	Summary table suffixes	Summary view suffixes
Weekly	_W	_WV
Monthly	_M	_MV
Quarterly	_Q	_QV
Yearly	_Y	_YV

Table 18 shows the expansion to summary columns of some of the most commonly used attribute types.

Table 18. 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 19 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 19. Primary key and warehouse changes for the Monitoring Agent for Linux OS

Attribute group (the attribute group name as it is displayed in the Tivoli Enterprise Portal)	Current primary keys	Removed primary keys	Warehoused
KLZ_CPU_Averages			Yes
KLZ_CPU	CPU_ID		Yes
KLZ_Disk_IO	Dev_Name		Yes
KLZ_Disk_Usage_Trends	Disk_Name		Yes
KLZ_Disk	Mount_Point Disk_Name		Yes
KLZ_IO_Ext	Device_Name		Yes
KLZ_NFS_Statistics	NFS_Version Location		Yes
KLZ_Network	Network_Interface_Name		Yes
KLZ_Process_User_Info	Process_ID		Yes
KLZ_Process	Process_ID		Yes
KLZ_RPC_Statistics			Yes
KLZ_Sockets_Detail	Socket_Inode		Yes
KLZ_Sockets_Status	Socket_Protocol		Yes
KLZ_Swap_Rate			Yes
KLZ_System_Statistics			Yes
KLZ_User_Login	Login_PID User_Name		Yes
KLZ_VM_Stats			Yes
Linux_All_Users	User_ID_64 User_ID		Yes
Linux_CPU_Averages			Yes
Linux_CPU_Config	CPU_ID		Yes
Linux_CPU	CPU_ID		Yes
Linux_Disk_IO	Dev_Name		Yes
Linux_Disk_Usage_Trends	Disk_Name		Yes
Linux_Disk	Mount_Point_U Disk_Name		Yes
Linux_File_Comparison	File_Name_2 File_Name_1		No
Linux_File_Information	File_Name_U Path_U		No
Linux_File_Pattern	File_Name		No
Linux_Group	Group_ID_64 Group_ID		Yes
Linux_Host_Availability	Target_Host		No
Linux_IO_Ext	Device_Name		Yes
Linux_IP_Address	IP_Address Network_Interface_Name		No

Table 19. Primary key and warehouse changes for the Monitoring Agent for Linux OS (continued)

Attribute group (the attribute group name as it is displayed in the Tivoli Enterprise Portal)	Current primary keys	Removed primary keys	Warehoused
Linux_Machine_Information			Yes
Linux_NFS_Statistics	NFS_Version Location		Yes
Linux_Network	Network_Interface_Name		Yes
Linux_OS_Config	OS_Name		Yes
Linux_Process_User_Info	Process_ID		Yes
Linux_Process	Process_ID		Yes
Linux_RPC_Statistics			Yes
Linux_Sockets_Detail	Socket_Inode		Yes
Linux_Sockets_Status	Socket_Protocol		Yes
Linux_Swap_Rate			Yes
Linux_System_Statistics			Yes
Linux_User_Login	User_Name_U Login_PID		Yes
Linux_VM_Stats			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 20 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 20. Scripts for affected attribute groups and summary tables for the Monitoring Agent for Linux OS

Attribute group or summary table	File	DB2	Oracle	MS SQL Server	Case 1	Case 2
Linux_All_Users	k1z_61migr_Linux_OS_Agent_Case-1.sql	X	X	X	X	
Linux_Group	k1z_61migr_Linux_OS_Agent_Case-1.sql	X	X	X	X	

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/sql1lib/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 load and alter 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 20 on page 232 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/sql1lib/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 on page 233. 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 20 on page 232 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/sql1lib/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. 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 21.

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 22 on page 240. 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 21. Overview of Distributed Monitoring migrated situations

Situation	IBM Tivoli Enterprise Console event class
LZ_USInodes*	Sentry2_0_inodes Sentry2_0_inodesused
LZ_USIUsPct*	Sentry2_0_inodesusedpct
LZ_USDkUPct*	Sentry2_0_diskusedpct
LZ_USDskAva*	Sentry2_0_diskavail
LZ_USDskUsd*	Sentry2_0_diskused
LZ_UStvDBSp*	Sentry2_0_tivdbspace
LZ_USDIORtK*	Sentry2_0_diskioratek
LZ_USRCPTmo*	Sentry2_0_rpctmout
LZ_USNtInEr*	Sentry2_0_netinerr
LZ_USNtInEX*	Sentry2_0_netinerrx

Table 21. Overview of Distributed Monitoring migrated situations (continued)

Situation	IBM Tivoli Enterprise Console event class
LZ_USNetIn*	Sentry2_0_netinerr
LZ_USNetInX*	Sentry2_0_netinx
LZ_USBadNFS*	Sentry2_0_badnfs
LZ_USBadNFS*	Sentry2_0_badnfs
LZ_USNetCol*	Sentry2_0_netcoll
LZ_USNCPct*	Sentry2_0_netcollpct
LZ_USNCPctX*	Sentry2_0_netcollpctx
LZ_USNetOEr*	Sentry2_0_netouterr
LZ_USNetOEX*	Sentry2_0_netouterrx
LZ_USNetOut*	Sentry2_0_netouterr
LZ_USNetOX*	Sentry2_0_netoutx
LZ_USBadRPC*	Sentry2_0_badrpc
LZ_USSwpAva*	Sentry2_0_swapavail
LZ_USCPUIdl*	Sentry2_0_cpuidle
LZ_USCPUSys*	Sentry2_0_cpusys
LZ_USCPUUsr*	Sentry2_0_cpuusr
LZ_USCPUSdu*	Sentry2_0_cpusdu
LZ_USCPUSpu*	Sentry2_0_cpuspu
LZ_USZombie*	Sentry2_0_zombies
LZ_USLdAv15*	Sentry2_0_loadavgfifteenm
LZ_USLdAv5*	Sentry2_0_loadavgonem
LZ_USLdAv1*	Sentry2_0_loadavgonem

Table 21. Overview of Distributed Monitoring migrated situations (continued)

Situation	IBM Tivoli Enterprise Console event class
LZ_USPgIns*	Sentry2_0_pageins
LZ_USPgOuts*	Sentry2_0_pageouts
LZ_USACPUbu*	Sentry2_0_avgcpubusy
LZ_UDskAva*	universal_diskavail
LZ_UDskUsd*	universal_diskused
LZ_UDskUPct*	universal_diskusedpct
LZ_UIndsFre*	universal_diskusedpct
LZ_UIndsUsd*	universal_diskusedpct
LZ_ULoadAvg*	universal_loadavg
LZ_UPageOut*	universal_pageouts
LZ_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 following table for attribute group to event classes and slots mapping information.

For example, if the situation is monitoring the No Password attribute from the All Users Group attribute group, the event class that is sent once the situation is triggered is ITM_Linux_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 KLZ_Base. The KLZ_Base event class can be used for generic rules processing for any event from the Monitoring Agent for Linux OS.

Table 22. Overview of attribute groups to event classes and slots

Attribute group	event classes and slots
Linux_User_Login	ITM_Linux_User_Login event class with these slots: <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • user_name: STRING • login_pid: INTEGER • login_pid_enum: STRING • line: STRING • login_time: STRING • idle_time: STRING • from_hostname: STRING • user_name_u: STRING • linux_vm_id: STRING
Linux_Disk	ITM_Linux_Disk event class with these slots: <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • disk_name: STRING • mount_point: STRING • size: INTEGER • size_enum: STRING • space_used: INTEGER • space_used_enum: STRING • space_available: INTEGER • space_available_enum: STRING • total_inodes: INTEGER • total_inodes_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 • mount_point_u: STRING • linux_vm_id: STRING • inodes_available_percent: INTEGER • inodes_available_percent_enum: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Disk_Usage_Trends	ITM_Linux_Disk_Usage_Trends event class with these slots: <ul style="list-style-type: none"> • system_name: STRING • timestamp: INTEGER • disk_name: STRING • space_used: INTEGER • space_used_enum: STRING • space_available: INTEGER • space_available_enum: STRING • disk_usage_rate: INTEGER • disk_usage_rate_enum: STRING • highwater_du_rate: INTEGER • highwater_du_rate_enum: STRING • highwater_time: STRING • disk_usage_moving_average: INTEGER • disk_usage_moving_average_enum: STRING • days_until_full_disk: INTEGER • days_until_full_disk_enum: STRING • days_full_disk_curr: INTEGER • days_full_disk_curr_enum: STRING • low_water_full_disk_curr: STRING • low_water_full_disk_curr_enum: STRING • days_full_disk_peak: INTEGER • days_full_disk_peak_enum: STRING • linux_vm_id: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Network	<p>ITM_Linux_Network event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: INTEGER • network_interface_name: STRING • interface_ip_address: STRING • interface_dns_name: STRING • interface_status: INTEGER • interface_status_enum: STRING • transmission_unit_maximum: INTEGER • transmission_unit_maximum_enum: STRING • kbytes_received_count: INTEGER • kbytes_received_count_enum: STRING • bytes_received_per_sec: INTEGER • bytes_received_per_sec_enum: STRING • kbytes_transmitted_count: INTEGER • kbytes_transmitted_count_enum: STRING • bytes_transmitted_per_sec: INTEGER • bytes_transmitted_per_sec_enum: STRING • packets_received_count: INTEGER • packets_received_count_enum: STRING • packets_received_per_sec: INTEGER • packets_received_per_sec_enum: STRING • input_errors: INTEGER • input_errors_enum: STRING • output_errors: INTEGER • output_errors_enum: STRING • packets_transmitted_per_sec: INTEGER • packets_transmitted_per_sec_enum: STRING • input_errors: INTEGER • input_errors_enum: STRING • output_errors: INTEGER • output_errors_enum: STRING • collisions: INTEGER • collisions_enum: STRING • collision_rate: INTEGER • collision_rate_enum: STRING • collision_percent: INTEGER

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Network (Continued)	<ul style="list-style-type: none"> • input_error_rate: INTEGER • input_error_rate_enum: STRING • output_error_rate: INTEGER • output_error_rate_enum: STRING • total_error_percent: INTEGER • input_packets_dropped: INTEGER • input_packets_dropped_enum: STRING • output_packets_dropped: INTEGER • output_packets_dropped_enum: STRING • input_fifo_buffer_overruns: INTEGER • input_fifo_buffer_overruns_enum: STRING • output_fifo_buffer_overruns: INTEGER • output_fifo_buffer_overruns_enum: STRING • packet_framing_errors: INTEGER • packet_framing_errors_enum: STRING • carrier_losses: INTEGER • carrier_losses_enum: STRING • linux_vm_id: STRING • input_error_percent: INTEGER • output_error_percent: INTEGER • device_type: INTEGER • device_type_enum: STRING • mac_address: STRING • mac_address_enum: STRING
User	<p>ITM_Linux_CPU event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • cpu_id: INTEGER • cpu_id_enum: STRING • user_cpu: REAL • user_nice_cpu: REAL • system_cpu: REAL • idle_cpu: REAL • busy_cpu: REAL • wait_io_cpu: REAL • user_sys_pct: INTEGER • steal_time_cpu: REAL • linux_vm_id: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_CPU_Averages	<p>ITM_Linux_CPU_Averages event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • days_to_cpu_upgrade: REAL • days_to_cpu_upgrade_enum: STRING • cpu_usage_current_average: REAL • cpu_usage_moving_average: REAL • user_nice_cpu_current_average: REAL • user_nice_cpu_moving_average: REAL • user_cpu_current_average: REAL • user_cpu_moving_average: REAL • system_cpu_current_average: REAL • system_cpu_moving_average: REAL • idle_cpu_current_average: REAL • idle_cpu_moving_average: REAL • wait_cpu_current_average: REAL • wait_cpu_moving_average: REAL • steal_cpu_current_average: REAL • steal_cpu_current_average_enum: STRING • steal_cpu_moving_average: REAL • steal_cpu_moving_average_enum: REAL • linux_vm_id: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Process	<p>ITM_Linux_Process event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: INTEGER • process_id: REAL • parent_process_id: INTEGER • process_command_name: STRING • state: INTEGER • state_enum: STRING • proc_system_cpu: REAL • proc_user_cpu: REAL • tot_proc_system_cpu: REAL • tot_proc_user_cpu: REAL • priority: INTEGER • nice: INTEGER • total_size_memory: INTEGER • total_size_memory_enum: STRING • resident_set_size: INTEGER • resident_set_size_enum: STRING • shared_memory: INTEGER • shared_memory_enum: STRING • text_resident_size: INTEGER • text_resident_size_enum: STRING • shared_lib_set_size: INTEGER • shared_lib_set_size_enum: STRING • data_set_size: INTEGER • data_set_size_enum: STRING • dirty_pages: INTEGER • dirty_pages_enum: STRING • vm_size: INTEGER • vm_size_enum: STRING • vm_lock: INTEGER • vm_lock_enum: STRING • vm_data: INTEGER • vm_data_enum: STRING • vm_stack: INTEGER • vm_stack_enum: STRING • vm_exe_size: INTEGER • vm_exe_size_enum: STRING • vm_lib_size: INTEGER • vm_lib_size_enum: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Process (continued)	<ul style="list-style-type: none"> • tot_minor_faults: INTEGER • tot_minor_faults_enum: STRING • tot_major_faults: INTEGER • tot_major_faults_enum: STRING • proc_cmd_line: STRING • proc_cmd_line_u: STRING • proc_cpu: INTEGER • proc_cpu_enum: STRING • linux_vm_id: STRING • user_sys_cpu_pct: INTEGER • process_command_name_u: STRING • total_busy_cpu_pct: REAL • busy_cpu_pct: REAL • vm_size_mb: REAL • vm_size_mb_enum: STRING • vm_lock_mb: REAL • vm_lock_mb_enum: STRING • vm_data_mb: REAL • vm_data_mb_enum: STRING • vm_stack_mb: REAL • vm_stack_mb_enum: STRING • vm_exe_size_mb: REAL • vm_exe_size_mb_enum: STRING • vm_lib_size_mb: REAL • vm_lib_size_mb_enum: STRING • threads: INTEGER • threads_enum: STRING • session_id: INTEGER • session_id_enum: STRING • proc_system_cpu_norm: REAL • proc_system_cpu_norm_enum: STRING • proc_user_cpu_norm: REAL • proc_user_cpu_norm_enum: STRING • proc_busy_cpu_norm: REAL • proc_busy_cpu_norm_enum: STRING • process_count: INTEGER • process_count_enum: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Process_User_Info	<p>ITM_Linux_Process_User_Info event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • process_id: INTEGER • process_id_enum: STRING • real_user_name: STRING • eff_user_name: STRING • saved_user_name: STRING • fs_user_name: STRING • real_group: STRING • eff_group: STRING • saved_group: STRING • file_sys_group: STRING • real_user_id: INTEGER • real_user_id_enum: STRING • eff_user_id: INTEGER • eff_user_id_enum: STRING • saved_user_id: INTEGER • saved_user_id_enum: STRING • fs_user_id: INTEGER • fs_user_id_enum: STRING • real_group_id: INTEGER • real_group_id_enum: STRING • eff_group_id: INTEGER • eff_group_id_enum: STRING • saved_group_id: INTEGER • saved_group_id_enum: STRING • file_sys_group_id: INTEGER • file_sys_group_id_enum: STRING • real_user_name_u: STRING • eff_user_name_u: STRING • saved_user_name_u: STRING • fs_user_name_u: STRING • real_group_u: STRING • eff_group_u: STRING • saved_group_u: STRING • file_sys_group_u: STRING • linux_vm_id: STRING • session_id: INTEGER • session_id_enum: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Process_User_Info (Continued)	<ul style="list-style-type: none"> • parent_process_id: INTEGER • parent_process_id_enum: STRING • state: INTEGER • state_enum: STRING • proc_cmd_line_u: STRING • process_command_name_u: STRING • vm_size_mb: REAL • vm_size_mb_enum: STRING • terminal_device: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_System_Statistics	<p>ITM_Linux_System_Statistics event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • ctxt_switches_per_sec: INTEGER • ctxt_switches_per_sec_enum: STRING • pct_change_ctxt_switches: REAL • processes_per_sec: INTEGER • pct_change_processes: REAL • pct_change_processes_enum: STRING • number_of_users: INTEGER • number_of_users_enum: STRING • system_load_1min: REAL • system_load_1min_enum: STRING • system_load_5min: REAL • system_load_5min_enum: STRING • system_load_15min: REAL • system_load_15min_enum: STRING • system_uptime: INTEGER • system_uptime_enum: STRING • linux_vm_id: STRING • pages_paged_in: INTEGER • pages_paged_in_enum: STRING • pages_paged_in_per_sec: REAL • pages_paged_in_per_sec_enum: STRING • pages_paged_out: INTEGER • pages_paged_out_enum: STRING • pages_paged_out_per_sec: REAL • pages_paged_out_per_sec_enum: STRING • pages_swapped_in: INTEGER • pages_swapped_in_enum: STRING • pages_swap_in_per_sec: REAL • pages_swap_in_per_sec_enum: STRING • pages_swapped_out: INTEGER • pages_swapped_out_enum: STRING • pages_swap_out_per_sec: REAL • pages_swap_out_per_sec_enum: STRING • page_faults_per_sec: INTEGER • page_faults_per_sec_enum: STRING • major_faults_per_sec: INTEGER • major_faults_per_sec_enum: STRING • total_number_processes: INTEGER • total_number_processes_enum: STRING • processes_zombie_count: INTEGER • processes_zombie_count_enum: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Swap_Rate	ITM_Linux_Swap_Rate event class with these slots: <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • moving_total_swap_space: INTEGER • moving_total_swap_space_enum: STRING • swap_space_used: INTEGER • swap_space_used_enum: STRING • swap_usage_rate: INTEGER • swap_usage_rate_enum: STRING • days_to_swap_space_full: INTEGER • days_to_swap_space_full_enum: STRING • peak_swap_space_used: INTEGER • peak_swap_space_used_enum: STRING • days_to_peak_space_full: INTEGER • days_to_peak_space_full_enum: STRING • low_free_memory: INTEGER • low_free_memory_enum: STRING • linux_vm_id: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_VM_Stats	ITM_Linux_VM_Stats event class with these slots: <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • total_swap_space: REAL • total_swap_space_enum: STRING • swap_space_used: REAL • swap_space_used_enum: STRING • swap_usage_free: REAL • swap_usage_free_enum: STRING • total_memory: REAL • total_memory_enum: STRING • memory_used: REAL • memory_used_enum: STRING • memory_free: REAL • memory_free_enum: STRING • shared_memory: REAL • shared_memory_enum: STRING • memory_in_buffers: REAL • memory_in_buffers_enum: STRING • memory_cached: REAL • memory_cached_enum: STRING • linux_vm_id: STRING • total_virtual_storage: REAL • total_virtual_storage_enum: STRING • used_virtual_storage: REAL • used_virtual_storage_enum: STRING • available_virtual_storage: REAL • available_virtual_storage_enum: STRING • virtual_storage_pct_avail: INTEGER • virtual_storage_pct_avail_enum: STRING • virtual_storage_pct_used: INTEGER • virtual_storage_pct_used_enum: STRING • real_memory_pct_used: INTEGER • real_memory_pct_used_enum: STRING • real_memory_pct_avail: INTEGER • real_memory_pct_avail_enum: STRING • swap_pct_used: INTEGER • swap_pct_used_enum: STRING • swap_pct_avail: INTEGER • swap_pct_avail_enum: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Sockets_Status	ITM_Linux_Sockets_Status event class with these slots: <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • socket_protocol: INTEGER • socket_protocol_enum: STRING • sockets_in_use: INTEGER • sockets_in_use_enum: STRING • highest_sockets_used: INTEGER • highest_sockets_used_enum: STRING • linux_vm_id: STRING
Linux_Sockets_Detail	ITM_Linux_Sockets_Detail event class with these slots: <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • socket_protocol: INTEGER • socket_protocol_enum: STRING • receive_queue: INTEGER • receive_queue_enum: STRING • send_queue: INTEGER • send_queue_enum: STRING • local_address: STRING • local_port: INTEGER • local_port_enum: STRING • local_service: STRING • foreign_address: STRING • socket_state: INTEGER • socket_state_enum: STRING • socket_uid: INTEGER • socket_uid_enum: STRING • socket_inode: INTEGER • socket_inode_enum: STRING • foreign_port: INTEGER • foreign_port_enum: STRING • socket_owner_name_u: STRING • linux_vm_id: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Disk_IO	ITM_Linux_Disk_IO event class with these slots: <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • transfers_per_sec: REAL • transfers_per_sec_enum: STRING • blk_rds_per_sec: REAL • blk_rds_per_sec_enum: STRING • blk_wrtn_per_sec: REAL • blk_wrtn_per_sec_enum: STRING • blks_read: INTEGER • blks_read_enum: STRING • blks_wrtn: INTEGER • blks_wrtn_enum: STRING • dev_major: INTEGER • dev_major_enum: STRING • dev_minor: INTEGER • dev_minor_enum: STRING • dev_name: STRING • linux_vm_id: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_IO_Ext	<p>ITM_Linux_IO_Ext event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • device_name: STRING • read_reqm_per_sec: REAL • read_reqm_per_sec_enum: STRING • write_reqm_per_sec: REAL • write_reqm_per_sec_enum: STRING • read_req_per_sec: REAL • read_req_per_sec_enum: STRING • write_req_per_sec: REAL • write_req_per_sec_enum: STRING • read_sect_per_sec: REAL • read_sect_per_sec_enum: STRING • write_sect_per_sec: REAL • write_sect_per_sec_enum: STRING • avg_req_size: REAL • avg_req_size_enum: STRING • avg_req_queue_length: REAL • avg_req_queue_length_enum: STRING • avg_wait_time: REAL • avg_wait_time_enum: STRING • avg_svc_time: REAL • avg_svc_time_enum: STRING • cpu_util: REAL • cpu_util_enum: STRING • linux_vm_id: STRING • disk_read_percent: REAL • disk_write_percent: REAL • read_bytes_per_sec: REAL • read_bytes_per_sec_enum: STRING • write_bytes_per_sec: REAL • write_bytes_per_sec_enum: STRING • transfers_bytes_per_sec: REAL • transfers_bytes_per_sec_enum: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_RPC_Statistics	<p>ITM_Linux_RPC_Statistics event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • rpc_server_total_calls: INTEGER • rpc_server_total_calls_enum: STRING • rpc_server_calls_rejected: INTEGER • rpc_server_calls_rejected_enum: STRING • rpc_server_packets_bad_auth: INTEGER • rpc_server_packets_bad_auth_enum: STRING • rpc_server_packets_bad_clt: INTEGER • rpc_server_packets_bad_clt_enum: STRING • rpc_server_packets_with_malformed_header: INTEGER • rpc_server_packets_with_malformed_header_enum: STRING • rpc_client_calls: INTEGER • rpc_client_calls_enum: STRING • rpc_client_calls_retransmitted: INTEGER • rpc_client_calls_retransmitted_enum: STRING • rpc_client_times_authentication_refreshed: INTEGER • rpc_client_times_authentication_refreshed_enum: STRING • linux_vm_id: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_NFS_Statistics	<p>ITM_Linux_NFS_Statistics event class with these slots:</p> <ul style="list-style-type: none"> • location: INTEGER • location_enum: STRING • nfs_version: INTEGER • nfs_version_enum: STRING • nfs_null_calls: INTEGER • nfs_null_calls_enum: STRING • nfs_null_call_percentage: INTEGER • nfs_null_call_percentage_enum: STRING • nfs_get_attribute_calls: INTEGER • nfs_get_attribute_calls_enum: STRING • nfs_get_attribute_calls_pct: INTEGER • nfs_get_attribute_calls_pct_enum: STRING • nfs_set_attribute_calls: INTEGER • nfs_set_attribute_calls_enum: STRING • nfs_set_attrib_calls_pct: INTEGER • nfs_set_attrib_calls_pct_enum: STRING • nfs_root_calls: INTEGER • nfs_root_calls_enum: STRING • nfs_root_calls_pct: INTEGER • nfs_root_calls_pct_enum: STRING • nfs_lookups: INTEGER • nfs_lookups_enum: STRING • nfs_lookups_pct: INTEGER • nfs_lookups_pct_enum: STRING • nfs_read_link_calls: INTEGER • nfs_read_link_calls_enum: STRING • nfs_read_link_pct: INTEGER • nfs_read_link_pct_enum: STRING • nfs_read_calls: INTEGER • nfs_read_calls_enum: STRING • nfs_read_calls_pct: INTEGER • nfs_read_calls_pct_enum: STRING • nfs_write_cache_calls: INTEGER • nfs_write_cache_calls_enum: STRING • nfs_write_cache_calls_pct: INTEGER • nfs_write_cache_calls_pct_enum: STRING • nfs_writes: INTEGER • nfs_writes_enum: STRING • nfs_writes_pct: INTEGER • nfs_writes_pct_enum: STRING • nfs_file_creates: INTEGER • nfs_file_creates_enum: STRING • nfs_file_creates_pct: INTEGER

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_NFS_Statistics (continued)	<ul style="list-style-type: none"> • nfs_file_creates_pct_enum: STRING • nfs_remove_file_calls: INTEGER • nfs_remove_file_calls_enum: STRING • nfs_remove_file_calls_pct: INTEGER • nfs_remove_file_calls_pct_enum: STRING • nfs_rename_file_calls: INTEGER • nfs_rename_file_calls_enum: STRING • rename_file_calls_pct: INTEGER • rename_file_calls_pct_enum: STRING • nfs_link_calls: INTEGER • nfs_link_calls_enum: STRING • link_calls_pct: INTEGER • link_calls_pct_enum: STRING • nfs_symbolic_link_calls: INTEGER • nfs_symbolic_link_calls_enum: STRING • symbolic_link_calls_pct: INTEGER • symbolic_link_calls_pct_enum: STRING • nfs_make_directory_calls: INTEGER • nfs_make_directory_calls_enum: STRING • nfs_make_directory_calls_pct: INTEGER • nfs_make_directory_calls_pct_enum: STRING • nfs_remove_directory_calls: INTEGER • nfs_remove_directory_calls_enum: STRING • remove_directory_calls_pct: INTEGER • remove_directory_calls_pct_enum: STRING • nfs_read_directory_calls: INTEGER • nfs_read_directory_calls_enum: STRING • read_directory_calls_pct: INTEGER • read_directory_calls_pct_enum: STRING • nfs_file_system_statistics_calls: INTEGER • nfs_file_system_statistics_calls_enum: STRING • file_system_statistics_calls_pct: INTEGER • file_system_statistics_calls_pct_enum: STRING • nfs_access: INTEGER • nfs_access_enum: STRING • access_pct: INTEGER • access_pct_enum: STRING • nfs_make_node_calls: INTEGER • nfs_make_node_calls_enum: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_NFS_Statistics (continued)	<ul style="list-style-type: none"> • make_node_calls_pct: INTEGER • make_node_calls_pct_enum: STRING • nfs_read_dir_calls_plus: INTEGER • nfs_read_dir_calls_plus_enum: STRING • read_dir_calls_plus_pct: INTEGER • read_dir_calls_plus_pct_enum: STRING • nfs_file_system_info: INTEGER • nfs_file_system_info_enum: STRING • file_system_info_pct: INTEGER • file_system_info_pct_enum: STRING • nfs_path_conf_calls: INTEGER • nfs_path_conf_calls_enum: STRING • path_conf_calls_pct: INTEGER • path_conf_calls_pct_enum: STRING • nfs_commit: INTEGER • nfs_commit_enum: STRING • nfs_commit_pct: INTEGER • nfs_commit_pct_enum: STRING • system_name: INTEGER • timestamp: STRING • linux_vm_id: STRING • nfs_total_calls: INTEGER • nfs_total_calls_enum: STRING
Linux_CPU_Config	<p>ITM_Linux_CPU_Config event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • cpu_id: INTEGER • vendor_id: STRING • cpu_family: INTEGER • cpu_family_enum: STRING • cpu_model: INTEGER • cpu_model_enum: STRING • model_name: STRING • clock_speed: REAL • clock_speed_enum: STRING • cache_size: INTEGER • cache_size_enum: STRING
Linux_OS_Config	<p>ITM_Linux_OS_Config event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • os_name: STRING • os_version: STRING • gcc_version: STRING • os_vendor: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_File_Information	<p>ITM_Linux_File_Information event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • path_u: STRING • file_name_u: STRING • size_mb: REAL • size_mb_enum: STRING • owner_u: STRING • group_u: STRING • last_changed_time: STRING • last_accessed_time: STRING • links: INTEGER • access: INTEGER • type: STRING • type_enum: STRING • link_name_u: STRING • mode: STRING • last_attr_chg_time: STRING • checksum_algorithm: INTEGER • checksum_algorithm_enum: STRING • checksum: STRING • file_content_changed: INTEGER • file_content_changed_enum: STRING
Linux_Host_Availability	<p>ITM_Linux_Host_Availability event class with these slots:</p> <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • target_host: STRING • host_availability: INTEGER • host_availability_enum: STRING • response_time: REAL • response_time_enum: STRING
Linux_File_Pattern	<p>ITM_Linux_File_Pattern event class with these slots:</p> <ul style="list-style-type: none"> • 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 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_File_Comparison	ITM_Linux_File_Comparison event class with these slots: <ul style="list-style-type: none"> • 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
Linux_All_Users	ITM_Linux_All_Users event class with these slots: <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • name: STRING • user_id: INTEGER • user_id_enum: INTEGER • password_null: INTEGER • password_null_enum: STRING • user_duplicated: INTEGER • user_duplicated_enum: STRING • user_sessions: INTEGER • user_sessions_enum: STRING
Linux_Group	ITM_Linux_Group event class with these slots: <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • group_name: STRING • group_id: INTEGER • group_id_enum: STRING • group_duplicated: INTEGER • group_duplicated_enum: STRING

Table 22. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Machine_Information	ITM_Linux_Machine_Information event class with these slots: <ul style="list-style-type: none"> • system_name: STRING • timestamp: STRING • klz_hostname: STRING • klz_hostname_enum: STRING • hardware_brand: STRING • hardware_brand_enum: STRING • hardware_model: STRING • hardware_model_enum: STRING • number_of_processors_online: INTEGER • number_of_processors_online_enum: STRING • number_of_processors_configured: INTEGER • number_of_processors_configured_enum: STRING • bios_version: STRING • bios_version_enum: STRING • bios_release: STRING • bios_release_enum: STRING • machine_serial: STRING • machine_serial_enum: STRING • system_board_uuid: STRING
Linux_IP_Address	ITM_Linux_IP_Address event class with these slots: <ul style="list-style-type: none"> • 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 C. Monitoring Agent for Linux OS data collection

In general, the Monitoring Agent for Linux OS 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 will not put a significant load on the system as it gathers the data.

KLZCPU attributes

The following table lists the mechanisms used to gather KLZCPU attributes.

Table 23. Mechanisms used to gather KLZCPU attributes

Attribute name	Collection method	Data source	Validation command or formula
ORIGINNODE	Short host name + ":LZ"	System file	N/A
TIMESTAMP	Current time	System file	N/A
CPUID	Read from /proc/stat	System file	mpstat -P ALL 30 ("CPU")
USRCPU	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample _interval * 10000	System file	mpstat -P ALL 30 ("%user")
USRNCPU	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample _interval * 10000	System file	mpstat -P ALL 30 ("%nice")
SYSCPU	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample _interval * 10000	System file	mpstat -P ALL 30 ("%system")
IDLECPU	10000 - BUSYCPU	System file	100 – Busy_CPU
BUSYCPU	USRCPU + USRNCPU + SYSCPU + WAITCPU	System file	User_CPU + User_Nice_CPU + System_CPU + Wait_IO_CPU + Steal_CPU_Percent
WAITCPU	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample _interval * 10000	System file	mpstat -P ALL 30 ("%iowait")
USRSYSCPU	((USRNCPU + USRCPU) * 100) / SYSCPU	System file	(User_Nice_CPU+User_CPU) / System_CPU
STEALCPU	Read from /proc/stat; samples_("N" - "N - 1") / Total_CPU_over_the_sample _interval * 10000	System file	mpstat -P ALL 30 ("%steal")

Note: All of the mpstat metrics are strictly dependent on the interval measure of the OS command and the API starting measure.

KLZCPUAVG attributes

The following table lists the mechanisms used to gather KLZCPUAVG attributes.

Table 24. Mechanisms used to gather KLZCPUAVG attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
DAYSCPU	Read from /proc/stat; total_moving_used_cpu / (previous_moving_idle - current_moving_idle); converted to days.
CPUCURAVG	USRNCURAVG + USRCURAVG + WAITCUR + SYSCPUCUR
CPUMOVAVG	Read from /proc/stat; metric_moving_average / moving_total_CPU_over_the_sample_interval * 10000; moving average of a metric is (previous_moving_average + samples_("N" - "N - 1")) / 2
USRNCURAVG	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample_interval * 10000
USRNMovCPU	Read from /proc/stat; metric_moving_average / moving_total_CPU_over_the_sample_interval * 10000; moving average of a metric is (previous_moving_average + samples_("N" - "N - 1")) / 2
USRCURAVG	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample_interval * 10000
USRMOVCPU	Read from /proc/stat; metric_moving_average / moving_total_CPU_over_the_sample_interval * 10000; moving average of a metric is (previous_moving_average + samples_("N" - "N - 1")) / 2
SYSCPUCUR	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample_interval * 10000
SYSCPUMOV	Read from /proc/stat; metric_moving_average / moving_total_CPU_over_the_sample_interval * 10000; moving average of a metric is (previous_moving_average + samples_("N" - "N - 1")) / 2
IDLECUR	10000 - CPUCURAVG
IDLEMOV	10000 - (USRNMovCPU + USRMOVCPU + WAITMOV+ SYSCPUMOV)
WAITCUR	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample_interval * 10000
WAITMOV	Read from /proc/stat; metric_moving_average / moving_total_CPU_over_the_sample_interval * 10000; moving average of a metric is (previous_moving_average + samples_("N" - "N - 1")) / 2

KLZDISK attributes

The following table lists the mechanisms used to gather KLZDISK attributes.

Table 25. Mechanisms used to gather KLZDISK attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
DSKNAME	getmntent API; struct mntent.mnt_fsname
MOUNTPT	getmntent API; struct mntent.mnt_dir
FSTYPE	statfs API; struct statfs elements: $f_blocks * (f_bsize / 1024) / 1024$
DSKSIZE	statfs API; struct statfs elements: $(f_blocks * (f_bsize / 1024)) / 1024$
DSKUSED	statfs API; struct statfs elements: $((f_blocks - f_bfree) * (f_bsize / 1024)) / 1024$
DSKUSEDPCT	$DSKUSED * 100.0 / (DSKUSED + DSKFREE)$
DSKFREE	statfs API; struct statfs elements: $((f_blocks - f_bfree) * (f_bsize / 1024)) / 1024$
DSKFREEPCT	$100 - DSKUSEDPCT$
INDSIZE	statfs API; struct statfs element: f_files
INDUSED	statfs API; struct statfs elements: $f_files - f_ffree$
INDFREE	statfs API; struct statfs element: f_ffree
INDFREEPCT	$100 - INDUSEDPCT$
INDUSEDPCT	$INODEUSED * 100.0 / f_files$
FSSTATUS	UP or DOWN (if statfs64() times out);

KLZDSKIO attributes

The following table lists the mechanisms used to gather KLZDSKIO attributes.

Table 26. Mechanisms used to gather KLZDSKIO attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
TPS	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); reads + writes; $samples_("N" - "N - 1") / sample_interval$
BLKRDSSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); sectors read; $samples_("N" - "N - 1") / sample_interval$
BLKWRTNSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); sectors written; $samples_("N" - "N - 1") / sample_interval$
BLKSRD	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); total sectors read
BLKSWRTN	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); total sectors written
DEVMAJOR	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel)

Table 26. Mechanisms used to gather KLZDSKIO attributes (continued)

Attribute name	Collection method
DEVMINOR	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel)
DKNAME	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel)

KLZDU attributes

The following table lists the mechanisms used to gather KLZDU attributes.

Table 27. Mechanisms used to gather KLZDU attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
DSKNAME	getmntent API; struct mntent.mnt_fsname
SPCUSED	statfs API; struct statfs elements: $((f_blocks - f_bfree) * (f_bsize / 1024)) / 1024$
SPCFREE	statfs API; struct statfs elements: $(f_bavail * (f_bsize / 1024)) / 1024$
DURATE	Calculated from "N" and "N - 1" samples of SPCUSED
HWDURATE	Larger of "N" and "N - 1" samples of DURATE
HWTIME	Timestamp associated with the HWDURATE sample
DUMVAVG	Average of all DURATE values
DAYSDSK	$(SPCAVAIL * 1024 * 1024) / (DUMVAVG * 24)$
DAYSCURR	$(SPCAVAIL * 1024 * 1024) / (DURATE * 24)$
LWCURR	Smaller of "N" and "N - 1" samples of DAYSCURR
DAYSPEAK	$(SPCAVAIL * 1024 * 1024) / (HWDURATE * 24)$

KLZIOEXT attributes

The following table lists the mechanisms used to gather KLZIOEXT attributes.

Table 28. Mechanisms used to gather KLZIOEXT attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
DKNAME	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel)
RDRQMSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); reads merged; $\text{samples}_{("N" - "N - 1")} / \text{sample_interval}$
WRTRQMSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); writes merged; $\text{samples}_{("N" - "N - 1")} / \text{sample_interval}$
RDRQSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); reads; $\text{samples}_{("N" - "N - 1")} / \text{sample_interval}$
WRTREQSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); writes; $\text{samples}_{("N" - "N - 1")} / \text{sample_interval}$

Table 28. Mechanisms used to gather KLZIOEXT attributes (continued)

Attribute name	Collection method
RDSECTSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); sectors read; samples_("N" - "N - 1") / sample_interval
WRSECTSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); sectors written; samples_("N" - "N - 1") / sample_interval
AVGRQSZ	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); samples_("N" - "N - 1"); (sectors_read + sectors_written) / (totals_reads + total_writes)
AVGRQQUSZ	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); I/O in progress; samples_("N" - "N - 1") / sample_interval
AVGWAITTM	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); samples_("N" - "N - 1"); (time_reading + time_writing) / (totals_reads + total_writes)
AVGSVCTM	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); samples_("N" - "N - 1"); time_spent_on_I/O / (totals_reads + total_writes)
IOUTIL	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); samples_("N" - "N - 1"); time_spent_on_I/O / monitoring_interval
IUTIL	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); samples_("N" - "N - 1"); IOUTIL / (total_reads / (totals_reads + total_writes))
OUTIL	IOUTIL - OUTIL
RDBYTESEC	RDSECTSEC converted to bytes
WRBYTESEC	WRSECTSEC converted to bytes
TOTBYTSEC	WRBYTESEC + RDBYTESEC

KLZLOGIN attributes

The following table lists the mechanisms used to gather KLZLOGIN attributes.

Table 29. Mechanisms used to gather KLZLOGIN attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
USERNAME	getutent API; struct utmp.ut_user
USRPID	getutent API; struct utmp.ut_pid
LINE	getutent API; struct utmp.ut_line
LOGINTIME	getutent API; struct utmp.ut_tv.tv_sec
IDLETIME	stat API on /dev/ut_line to get last access time & subtract from current time
FROMHOST	getutent API; struct utmp.ut_host

KLZLPAR attributes

The following table lists the mechanisms used to gather KLZLPAR attributes.

Table 30. Mechanisms used to gather KLZLPAR attributes

Attribute name	Collection method
CAPWEIGHT	Read from /proc/ppc64/lparcfg
CAPPED	Read from /proc/ppc64/lparcfg
ENTITLEMENT	Read from /proc/ppc64/lparcfg
ENTITLUSED	PHCPUUSED*10000/ ENTITLEMENT
LPARID	Read from /proc/ppc64/lparcfg
MAXCAPUSED	PHCPUUSED*100/ NVIRTCPU (if CAPPED = 0); == ENTITLUSED (if CAPPED = 1)
NVIRTCPU	Read from /proc/ppc64/lparcfg
PHCPUUSED	Read from /proc/ppc64/lparcfg; samples_("N" - "N - 1") *100 / (sample_interval * timebase); timebase read from /proc/cpuinfo
SHARED	Read from /proc/ppc64/lparcfg
SHPOOLIDLE	Read from /proc/ppc64/lparcfg; samples ("N" - "N - 1") *100 / (sample_interval * timebase); timebase read from /proc/cpuinfo
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
STEALTIME	Read aggregate CPU from /proc/stat; samples_("N" - "N - 1") / Total_CPU_over_the_sample_interval * 10000

KLZNET attributes

The following table lists the mechanisms used to gather KLZNET attributes.

Table 31. Mechanisms used to gather KLZNET attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
FNAME	Read from /proc/net/dev
FIPADDR	socket, ioctl & inet_ntoa APIs
FSTATUS	socket & ioctl APIs
FMTU	socket & ioctl APIs
FIKBYTES	Read from /proc/net/dev & divided by 1024
RECBPS	Read from /proc/net/dev; samples_("N" - "N - 1") / sample_interval
FOKBYTES	Read from /proc/net/dev & divided by 1024
TRANSBPS	Read from /proc/net/dev; samples_("N" - "N - 1") / sample_interval
FIFRAMES	Read from /proc/net/dev
RPACKPS	Read from /proc/net/dev; samples_("N" - "N - 1") / sample_interval
FOFRAMES	Read from /proc/net/dev

Table 31. Mechanisms used to gather KLZNET attributes (continued)

Attribute name	Collection method
TPACKPS	Read from /proc/net/dev; samples_("N" - "N - 1") / sample_interval
FIERRORS	Read from /proc/net/dev
FOERRORS	Read from /proc/net/dev
FCOLLSNS	Read from /proc/net/dev
FCOLLSNRT	Read from /proc/net/dev; samples_("N" - "N - 1") * 60 / sample_interval
FCOLLSPCT	Read from /proc/net/dev; for this sample period: (collisions / (frames sent + frames rcved)) * 100
FIERRORT	Read from /proc/net/dev; samples_("N" - "N - 1") * 60 / sample_interval
FOERRORT	Read from /proc/net/dev; samples_("N" - "N - 1") * 60 / sample_interval
FIOERRPCT	Read from /proc/net/dev; for this sample period: (input_errors + output_errors) / (frames_sent + frames_rcved) * 100
FIDROP	Read 5th value from /proc/net/dev
FODROP	Read 11th value from /proc/net/dev
FIFOINOV	Read from /proc/net/dev
FIPKTFRAM	Read from /proc/net/dev
FCARRIER	Read from /proc/net/dev
FIERRPCT	FIOERRPCT * (FIERRORT / (FIERRORT + FOERRORT))
FOERRPCT	FIOERRPCT - FIERRPCT
DEVTYPE	socket & ioctl APIs
MACADDRESS	socket & ioctl APIs

KLZNFS attributes

The following table lists the mechanisms used to gather KLZNFS attributes.

Table 32. Mechanisms used to gather KLZNFS attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
LOCORIG	Generated client/server indicator
NFSVER	Generated version indicator
NFSNULL	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
NULLPCT	NFSNULL * 100 / NFSTOT
NFSGETATT	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
GETATTPCT	NFSGETADD * 100 / NFSTOT
NFSSETATT	
SETATTPCT	NFSSETATT * 100 / NFSTOT

Table 32. Mechanisms used to gather KLZNFS attributes (continued)

Attribute name	Collection method
NFSROOT	
NFSROOTPCT	NFSROOT * 100 / NFSTOT
NFSLOOKUP	
LOOKUPPCT	NFSLOOKUP * 100 / NFSTOT
NFSRDLINK	
RDLINKPCT	NFSRDLINK * 100 / NFSTOT
NFSREAD	
READPCT	NFSREAD * 100 / NFSTOT
NFSWRCACH	
WRCACHPCT	NFSWRCACH * 100 / NFSTOT
NFSWRITES	
NFSWRTPCT	NFSWRITES * 100 / NFSTOT
NFSCREATES	
CREATESPCT	NFSCREATES * 100 / NFSTOT
NFSREMOVE	
REMOVEPCT	NFSREMOVE * 100 / NFSTOT
NFSRENAME	
RENMPCT	NFSRENAME * 100 / NFSTOT
NFSLINK	
LINKPCT	NFSLINK * 100 / NFSTOT
NFSSYMLNK	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
SYMLNKPCT	NFSSYMLNK * 100 / NFSTOT
NFSMKDIR	
MKDIRPCT	NFSMKDIR * 100 / NFSTOT
NSRMDIR	
RMDIRPCT	NFSRMDIR * 100 / NFSTOT
NFSRDDIR	
RDDIRPCT	NFSRDDIR * 100 / NFSTOT
NFSFSSTAT	
FSSTATPCT	NFSFSSTAT * 100 / NFSTOT
NFSACCESS	
ACCSPCT	NFSACCESS * 100 / NFSTOT
NFSMKNOD	
MKNODPCT	NFSMKNOD * 100 / NFSTOT
RDDIRPLUS	
RDIRPLSPCT	RDDIRPLUS * 100 / NFSTOT
NFSFSINFO	
FSINFOPCT	NFSFSINFO * 100 / NFSTOT
NFSPTHCONF	

Table 32. Mechanisms used to gather KLZNFS attributes (continued)

Attribute name	Collection method
PTHCONF PCT	NFSPHCONF * 100 / NFSTOT
NFSCOMMIT	
COMMITPCT	NFSCOMMIT * 100 / NFSTOT
NFSTOT	NFSNULL + NFSGETATT + NFSSETATT + NFSROOT + NFSLOOKUP + NFSRDLINK + NFSREAD + NFSWRACACH + NFSWRITES + NFScreates + NFSREMOVE + NFSRENAME + NFSLINK + NFSSYMLNK + NFSMKDIR + NFSRMDIR + NFSRDDIR + NFSFSSTAT

KLZPROC attributes

The following table lists the mechanisms used to gather KLZPROC attributes.

Table 33. Mechanisms used to gather KLZPROC attributes

Attribute name	Collection method	Data source	Validation command or formula
ORIGINNODE	Short host name + ":LZ"	System file	N/A
TIMESTAMP	Current time	System file	N/A
PID	Read from /proc; PID is the subdirectory name	System file	ps -ef ("PID")
PPID	Read from /proc/PID/stat	System file	ps -p Process_ID -o ppid
CMD	Read from /proc/PID/stat	System file	ps -p Process_ID ("CMD")
CMDLINE	Read from /proc/PID/cmdline	System file	ps -p Process_ID -o cmd
STATE	Read from /proc/PID/stat	System file	ps -p Process_ID -o state
PSYSCPU	Read from /proc/PID/stat; converted from jiffies	System file	((cat /proc/Process_ID/stat (fifteenth field) *100) / getconf CLK_TCK) / (cat /proc/uptime (first field) - (cat /proc/Process_ID/stat (twentysecond field) / getconf CLK_TCK))
PUSRCPU	Read from /proc/PID/stat; converted from jiffies	System file	((cat /proc/Process_ID/stat (fourteenth field) *100) / getconf CLK_TCK) / (cat /proc/uptime (first field) - (cat /proc/Process_ID/stat (twentysecond field) / getconf CLK_TCK))

Table 33. Mechanisms used to gather KLZPROC attributes (continued)

Attribute name	Collection method	Data source	Validation command or formula
TSYSCPU	Read from /proc/PID/stat; converted from jiffies	System file	((cat /proc/Process_ID/stat (seventeenth field) *100) / getconf CLK_TCK) / (cat /proc/uptime (first field) - (cat /proc/Process_ID/stat (twentysecond field) / getconf CLK_TCK))
TUSRCPU	Read from /proc/PID/stat; converted from jiffies	System file	((cat /proc/Process_ID/stat (sixteenth field) *100) / getconf CLK_TCK) / (cat /proc/uptime (first field) - (cat /proc/Process_ID/stat (twentysecond field) / getconf CLK_TCK))
INTPRI	Read from /proc/PID/stat	System file	cat /proc/Process_ID/stat (seventeenth field) OR ps -p Process_ID -o priority
NICE	Read from /proc/PID/stat	System file	cat /proc/Process_ID/stat (eighteenth field) OR ps -p Process_ID -o nice
SIZE	Read from /proc/PID/statm	System file	cat /proc/Process_ID/stat (twenty-second field) OR (ps -p Process_ID -o vsize) / 4
RSS	Read from /proc/PID/statm	System file	cat /proc/Process_ID/stat (twenty-third field) OR (ps -p Process_ID -o rss) / 4
SHAREMEM	Read from /proc/PID/statm	System file	cat /proc/Process_ID/statm (third field)
TRS	Read from /proc/PID/statm	System file	cat /proc/Process_ID/statm (fifth field) OR (ps -p Process_ID -o tsiz) / 4
LRS	Read from /proc/PID/statm	System file	cat /proc/Process_ID/statm (fifth field)
DRS	Read from /proc/PID/statm	System file	cat /proc/Process_ID/statm (sixth field)
DIRTPG	Read from /proc/PID/statm	System file	cat /proc/Process_ID/statm (seventh field)
VMSIZE	Read from /proc/PID/status	System file	cat /proc/Process_ID/status ("VmSize")
VMLOCK	Read from /proc/PID/status	System file	cat /proc/Process_ID/status ("VmLck")
VMDATA	Read from /proc/PID/status	System file	cat /proc/Process_ID/status ("VmData")
VMSTACK	Read from /proc/PID/status	System file	cat /proc/Process_ID/status ("VmStk")
VMEXESZ	Read from /proc/PID/status	System file	cat /proc/Process_ID/status ("VmExe")
VMLIBSZ	Read from /proc/PID/status	System file	cat /proc/Process_ID/status ("VmLib")

Table 33. Mechanisms used to gather KLZPROC attributes (continued)

Attribute name	Collection method	Data source	Validation command or formula
CMINFLT	Read from /proc/PID/stat	System file	cat /proc/Process_ID/stat (eleventh field)
CMAJFLT	Read from /proc/PID/stat	System file	cat /proc/Process_ID/stat (thirteenth field)
CPUAFF	Read from /proc/PID/stat	System file	cat /proc/Process_ID/stat (thirty-third field) OR ps -p Process_ID -o psr
USRSYSCPU	(TUSRCPU / TSYSCPU) * 100	System file	Tot_Proc_User_CPU * 100 / Tot_Proc_System_CPU
TBUSYCPU	TSYSCPU + TUSRCPU	System file	ps -p Process_ID S -o pcpu OR (Cumulative Process User CPU (Percent)+Cumulative Process System CPU (Percent))
BUSYCPU	PSYSCPU + PUSRCPU	System file	ps -p Process_ID -o pcpu OR (Process User CPU (Percent)+ Process System CPU (Percent))
CPUSECONDS	Read user + system CPU time from /proc/PID/stat; converted from jiffies to seconds	System file	ps -p Process_ID -o cputime (→ seconds)
TOTALTIME	Read user+system CPU time from /proc/PID/stat; converted in days,hours,minutes,seconds format	System file	ps -p Process_ID -o cputime (→ hh mm ss) OR((Process User CPU (Time)+Process System CPU (Time))
UPROCFILT	the regular expression that matched; empty otherwise	System file	metric not collected for this platform
CPUPERCENT	Read user+system CPU time from /proc/PID/stat; converted from jiffies; samples ("N" - "N - 1") *10000 / (sample_interval * number of CPUs)	System file	top -p Process_ID -d 60 ("%CPU") / (cat /proc/cpuinfo grep -c processor)
SYSTEMTIM	Read system CPU time from /proc/PID/stat; converted in days,hours,minutes,seconds format	System file	cat /proc/Process_ID/stat (fifteenth field) / getconf CLK_TCK (→ hh mm ss)
USERTIME	Read user CPU time from /proc/PID/stat; converted in days,hours,minutes,seconds format	System file	cat /proc/Process_ID/stat (fourteenth field) / getconf CLK_TCK (→ hh mm ss)
VMSIZEMB	Read from /proc/PID/status; converted to MB	System file	cat /proc/Process_ID/status ("VmSize") / 1024

Table 33. Mechanisms used to gather KLZPROC attributes (continued)

Attribute name	Collection method	Data source	Validation command or formula
VMLOCKMB	Read from /proc/PID/status; converted to MB	System file	cat /proc/Process_ID/status ("VmLck") / 1024
VMDATAMB	Read from /proc/PID/status; converted to MB	System file	cat /proc/Process_ID/status ("VmData") / 1024
VMSTACKMB	Read from /proc/PID/status; converted to MB	System file	cat /proc/Process_ID/status ("VmStk") / 1024
VMEXESZMB	Read from /proc/PID/status; converted to MB	System file	cat /proc/Process_ID/status ("VmExe") / 1024
VMLIBSZMB	Read from /proc/PID/status; converted to MB	System file	cat /proc/Process_ID/status ("VmLib") / 1024
PROCTHRD	Read from /proc/PID/status	System file	(ps -p Process_ID -T grep -c Process ID) OR (cat /proc/Process_ID/status ("Threads"))
SESSIONID	Read from /proc/PID/stat	System file	ps -p Process_ID -o sid OR cat /proc/Process_ID/stat (sixth field)
PSYSNORM	Read from /proc/PID/stat; converted from jiffies	System file	((cat /proc/Process ID/stat (fourteenth field)(t-1) - cat /proc/Process ID/stat (fourteenth field)(t))*10000/hertz_)/((cat /proc/Process ID/stat (fifteenth field)(t-1) - cat /proc/Process ID/stat (fifteenth field)(t))*sysconf (_SC_NPROCESSORS_ONLN));
PUSRNORM	Process user-mode time read from /prod/PID/stat; Nbr of CPUs read from sysconf API; (current_user_mode - previous_user_mode) / (elapsed_time * nbr_of_cpus)	System file	((cat /proc/Process ID/stat (fifteenth field)(t-1) - cat /proc/Process ID/stat (fifteenth field)(t))*10000/hertz_)/((cat /proc/Process ID/stat (fifteenth field)(t-1) - cat /proc/Process ID/stat (fifteenth field)(t))*sysconf (_SC_NPROCESSORS_ONLN));
PBUSYNORM	Process kernel-mode time read from /prod/PID/stat; Nbr of CPUs read from sysconf API; (current_kernel_mode - previous_kernel_mode) / (elapsed_time * nbr_of_cpus)	System file	Process Short Term Avg System CPU (Percent)+Process Short Term Avg User CPU (Percent)

Table 33. Mechanisms used to gather KLZPROC attributes (continued)

Attribute name	Collection method	Data source	Validation command or formula
PROCCOUNT	Generated; count of processes with same CMDLINE	System file	(ps -eo cmd grep -c Command Line) -1

KLZPUSR attributes

The following table lists the mechanisms used to gather KLZPUSR attributes.

Table 34. Mechanisms used to gather KLZPUSR attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
PID	Read from /proc; PID is the subdirectory name
RUSERID	Read from /proc/PID/status
EUSERID	Read from /proc/PID/status
SUSERID	Read from /proc/PID/status
FSUSERID	Read from /proc/PID/status
RGRPID	Read from /proc/PID/status
EFFGRPID	Read from /proc/PID/status
SGRPID	Read from /proc/PID/status
FSGRPID	Read from /proc/PID/status
RUSER	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
EUSER	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
SUSER	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
FSUSER	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
RGRP	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
EGRP	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
SGRP	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
FSGRP	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
SESSIONID	Read from /proc/PID/stat
PPID	Read from /proc/PID/stat
STATE	Read from /proc/PID/stat
CMD	Read from /proc/PID/stat
CMDLINE	Read from /proc/PID/cmdline
VMSIZEMB	Read from /proc/PID/status; converted to MB

Table 34. Mechanisms used to gather KLZPUSR attributes (continued)

Attribute name	Collection method
TTY	Read from /proc/PID/stat; converted to string by internal method

KLZRPC attributes

The following table lists the mechanisms used to gather KLZRPC attributes.

Table 35. Mechanisms used to gather KLZRPC attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
RSCALLS	Read from /proc/net/rpc/nfsd
RSBADCALL	Read from /proc/net/rpc/nfsd
RSBADAUTH	Read from /proc/net/rpc/nfsd
RSBADCLT	Read from /proc/net/rpc/nfsd
RSXDRCALL	Read from /proc/net/rpc/nfsd
RCCALLS	Read from /proc/net/rpc/nfs
RCRETRAN	Read from /proc/net/rpc/nfs
RCAREF	Read from /proc/net/rpc/nfs

KLZSOCKD attributes

The following table lists the mechanisms used to gather KLZSOCKD attributes.

Table 36. Mechanisms used to gather KLZSOCKD attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
SCKPROTO	Generated TCP/UDP indicator
RECVQ	Read from /proc/net/tcp or /proc/net/udp
SENDQ	Read from /proc/net/tcp or /proc/net/udp
LOCLADDR	Read from /proc/net/tcp or /proc/net/udp
LOCLPORT	Read from /proc/net/tcp or /proc/net/udp
LOCLSVC	Read from /proc/net/tcp or /proc/net/udp & getservbyport API; struct servent.s_name
FORNADDR	Read from /proc/net/tcp or /proc/net/udp
STATE	Read from /proc/net/tcp or /proc/net/udp
SOCKUID	Read from /proc/net/tcp or /proc/net/udp
SCKINOD	Read from /proc/net/tcp or /proc/net/udp
REMOTPORT	Read from /proc/net/tcp or /proc/net/udp
RUSER	Read from /proc/net/tcp or /proc/net/udp & getpwuid API; struct passwd.pw_name

KLZSOCKS attributes

The following table lists the mechanisms used to gather KLZSOCKS attributes.

Table 37. Mechanisms used to gather KLZSOCKS attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
SCKPROTO	Read from /proc/net/sockstat
SCKINUSE	Read from /proc/net/sockstat
SCKHWUSED	Read from /proc/net/sockstat

KLZSWPRT attributes

The following table lists the mechanisms used to gather KLZSWPRT attributes.

Table 38. Mechanisms used to gather KLZSWPRT attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
MOVSWPTOT	Read from /proc/meminfo; (last MOVSWAPTOT + SwapTotal) / 2
SWAPUSED	Read from /proc/meminfo; (last SWAPUSED + (SwapTotal - SwapFree)) / 2
SWPRATE	Read from /proc/meminfo; (last SWAPRATE + ((SwapTotal - SwapFree) - previous_SWAPUSED)) / 2
SWAPDAYS	Read from /proc/meminfo; SwapTotal / (24 * SWAPRATE)
PKSWPUSD	Read from /proc/meminfo; larger of last two (SwapTotal - SwapFree)
MINSWPDAYS	Read from /proc/meminfo; smaller of last two SWAPDAYS
LOWMEM	Read from /proc/meminfo; LowFree

KLZSYS attributes

The following table lists the mechanisms used to gather KLZSYS attributes.

Table 39. Mechanisms used to gather KLZSYS attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
CSWSEC	Read from /proc/PID/stat; samples_("N" - "N - 1") / sample_interval
RATECSW	Read from /proc/PID/stat; ((current_CSWSEC - previous_CSWSEC) / previous_CSWSEC) * 100
PROCSEC	Read from /proc/PID/stat; samples_("N" - "N - 1") / sample_interval
RATEPROC	Read from /proc/PID/stat; ((current_PROCSEC - previous_PROCSEC) / previous_PROCSEC) * 100

Table 39. Mechanisms used to gather KLZSYS attributes (continued)

Attribute name	Collection method
CURUSRS	getutent API; count of entries in utmp database
LOAD1MIN	Read from /proc/loadavg
LOAD5MIN	Read from /proc/loadavg
LOAD15MIN	Read from /proc/loadavg
SYSUPTIME	Read from /proc/uptime
PGPGIN	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel)
PGPGINPS	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel); samples_("N" - "N - 1") / sample_interval * 100
PGPGOUT	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel)
PGPGOUTPS	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel); samples_("N" - "N - 1") / sample_interval * 100
PGSWAPIN	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel)
SWAPINPS	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel); samples_("N" - "N - 1") / sample_interval * 100
PGSWAPOUT	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel)
SWAPOUTPS	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel); samples_("N" - "N - 1") / sample_interval * 100
PGFLTPTS	Read from /proc/vmstat (2.6 kernel) samples_("N" - "N - 1") / sample_interval * 100; N/A on 2.4 kernel
MAJFLTPTS	Read from /proc/vmstat (2.6 kernel) samples_("N" - "N - 1") / sample_interval * 100; N/A on 2.4 kernel
TOTPROCS	Count process subdirs in /proc
ZOMBCNT	Count process subdirs in /proc in zombie state

KLZVM attributes

The following table lists the mechanisms used to gather KLZVM attributes.

Table 40. Mechanisms used to gather KLZVM attributes

Attribute name	Collection method	Data source	Validation command or formula
ORIGINNODE	Short host name + ":LZ"	System file	N/A
TIMESTAMP	Current time	System file	N/A
SWPTOT	Read from /proc/meminfo; (SwapTotal / 1024) * 100	System file	free ("Swap: total") / 1024
SWPUSED	Read from /proc/meminfo; ((SwapTotal - SwapFree) / 1024) * 100	System file	free ("Swap: used") / 1024
SWPUSEDPCT	Read from /proc/meminfo; ((SwapTotal - SwapFree) / SwapTotal) * 100	System file	Swap_Space_Used / Total_Swap_Space * 100
SWPFREE	Read from /proc/meminfo; (SwapFree / 1024) * 100	System file	free ("Swap: free") / 1024

Table 40. Mechanisms used to gather KLZVM attributes (continued)

Attribute name	Collection method	Data source	Validation command or formula
SWPFREEPCT	100 - SWPUSEDPCT	System file	100 - Swap_Used_Pct
MEMTOT	Read from /proc/meminfo; (MemTotal / 1024) * 100	System file	free ("Mem: total") / 1024
MEMUSED	Read from /proc/meminfo; ((MemTotal - MemFree) / 1024) * 100	System file	free ("Mem: used") / 1024
MEMUSEDPCT	Read from /proc/meminfo; ((MemTotal - MemFree) / MemTotal) * 100	System file	(Memory_Used / Total_Memory) * 100
MEMFREE	Read from /proc/meminfo; (MemFree / 1024) * 100	System file	free ("Mem: free") / 1024
MEMFREEPCT	100 - MEMUSEDPCT	System file	100 - Memory_Used_Pct
MEMSHARED	Read from /proc/meminfo; (MemShared / 1024) * 100	System file	free ("Mem: shared") / 1024
MEMBUFF	Read from /proc/meminfo; (Buffers / 1024) * 100	System file	free ("Mem: buffers") / 1024
MEMCACHE	Read from /proc/meminfo; (Cache / 1024) * 100	System file	free ("Mem: cached") / 1024
VSTOT	MEMTOT + SWPTOT	System file	Total_Memory + Total_Swap_Space
VSUSED	SWPUSEDPCT + MEMUSED	System file	Memory_Used + Swap_Space_Used
VSUSEDPCT	100 - VSFREEPCT	System file	100 - Free_Virtual_Storage_Pct
VSFREE	VSTOT - VSUSED	System file	Total_Virtual_Storage - Used_Virtual_Storage
VSFREEPCT	(VSFREE / VSTOT) * 100	System file	Free_Virtual_Storage / Total_Virtual_Storage * 100
MEMCACHPCT	Read cached and total from /proc/meminfo; cached * 100/total	System file	Memory_Cached / Total_Memory * 100
NETMEMUSED	MEMUSED - MEMCACHE - MEMBUFF	System file	free ("-/+ buffers/cache: used") / 1024
NETMEMUPCT	NETMEMUSED / MEMTOT	System file	Net_Memory_Used / Total_Memory * 100
TOTMEMFREE	MEMFREE + MEMCACHE + MEMBUFF	System file	free ("-/+ buffers/cache: free") / 1024
TOTMEMFPCT	TOTMEMFREE / MEMTOT	System file	Total_Memory_Free / Total_Memory * 100

LNXALLUSR attributes

The following table lists the mechanisms used to gather LNXALLUSR attributes.

Table 41. Mechanisms used to gather LNXALLUSR attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
USERNAME	getpwent API; struct passwd.pw_name
UID	getpwent API; struct passwd.pw_uid
PWNULL	getspnam API; generated; true if struct spwd.sp_pwdp is empty
USERDUP	Generated; true if duplicate USERNAME or UID is detected
USERSES	getpwent & getutxent APIs; generated; matches of struct paswd.pw_name & struct utmpx.ut_user fields
UID64	getpwent API; struct passwd.pw_uid

LNXCPU attributes

The following table lists the mechanisms used to gather LNXCPU attributes.

Table 42. Mechanisms used to gather LNXCPU attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
CPUID	Read from /proc/stat
USRCPU	Read from /proc/stat; $\text{samples_}("N" - "N - 1") / \text{total_CPU_over_the_sample_interval} * 10000$
USRNCPU	Read from /proc/stat; $\text{samples_}("N" - "N - 1") / \text{total_CPU_over_the_sample_interval} * 10000$
SYSCPU	Read from /proc/stat; $\text{samples_}("N" - "N - 1") / \text{total_CPU_over_the_sample_interval} * 10000$
IDLECPU	$10000 - \text{BUSYCPU}$
BUSYCPU	$\text{USRCPU} + \text{USRNCPU} + \text{SYSCPU} + \text{WAITCPU}$
WAITCPU	Read from /proc/stat; $\text{samples_}("N" - "N - 1") / \text{total_CPU_over_the_sample_interval} * 10000$
USRSYSCPU	$((\text{USRNCPU} + \text{USRCPU}) * 100) / \text{SYSCPU}$

LNXCPUAVG attributes

The following table lists the mechanisms used to gather LNXCPUAVG attributes.

Table 43. Mechanisms used to gather LNXCPUAVG attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
DAYSCPU	Read from /proc/stat; $\text{total_moving_used_cpu} / (\text{previous_moving_idle} - \text{current_moving_idle})$; converted to days.

Table 43. Mechanisms used to gather LNXCPUAVG attributes (continued)

Attribute name	Collection method
CPUCURAVG	USRNCURAVG + USRCURAVG + WAITCUR + SYSCPUCUR
CPUMOVAVG	Read from /proc/stat; metric_moving_average / moving_total_CPU_over_the_sample_interval * 10000; moving average of a metric is (previous_moving_average + samples_("N" - "N - 1")) / 2
USRNCURAVG	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample_interval * 10000
USRNMOVCPU	Read from /proc/stat; metric_moving_average / moving_total_CPU_over_the_sample_interval * 10000; moving average of a metric is (previous_moving_average + samples_("N" - "N - 1")) / 2
USRCURAVG	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample_interval * 10000
USRMOVCPU	Read from /proc/stat; metric_moving_average / moving_total_CPU_over_the_sample_interval * 10000; moving average of a metric is (previous_moving_average + samples_("N" - "N - 1")) / 2
SYSCPUCUR	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample_interval * 10000
SYSCPUMOV	Read from /proc/stat; metric_moving_average / moving_total_CPU_over_the_sample_interval * 10000; moving average of a metric is (previous_moving_average + samples_("N" - "N - 1")) / 2
IDLECUR	10000 - CPUCURAVG
IDLEMOV	10000 - (USRNMOVCPU + USRMOVCPU + WAITMOV+ SYSCPUMOV)
WAITCUR	Read from /proc/stat; samples_("N" - "N - 1") / total_CPU_over_the_sample_interval * 10000
WAITMOV	Read from /proc/stat; metric_moving_average / moving_total_CPU_over_the_sample_interval * 10000; moving average of a metric is (previous_moving_average + samples_("N" - "N - 1")) / 2

LNXCPUCON attributes

The following table lists the mechanisms used to gather LNXCPUCON attributes.

Table 44. Mechanisms used to gather LNXCPUCON attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
CPUID	Read from /proc/cpuinfo
VENDID	Read from /proc/cpuinfo
CPUFAMILY	Read from /proc/cpuinfo
CPUMODEL	Read from /proc/cpuinfo
MODELNM	Read from /proc/cpuinfo
CPUCLK	Read from /proc/cpuinfo

Table 44. Mechanisms used to gather LNXCPUCON attributes (continued)

Attribute name	Collection method
CACHESZ	Read from /proc/cpuinfo

LNXDISK attributes

The following table lists the mechanisms used to gather LNXDISK attributes.

Table 45. Mechanisms used to gather LNXDISK attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
DSKNAME	getmntent API; struct mntent.mnt_fsname
MOUNTPT	getmntent API; struct mntent.mnt_dir
DSKSIZE	statfs API; struct statfs elements: $f_blocks * (f_bsize / 1024) / 1024$
SPCUSED	statfs API; struct statfs elements: $((f_blocks - f_bfree) * (f_bsize / 1024)) / 1024$
SPCAVAIL	statfs API; struct statfs elements: $(f_bavail * (f_bsize / 1024)) / 1024$
INODESIZE	statfs API; struct statfs element: f_files
INODEUSED	statfs API; struct statfs elements: $f_files - f_ffree$
INODEFREE	statfs API; struct statfs element: f_ffree
PCTSPCUSED	$SPCUSED * 100.0 / (SPCUSED + SPCAVAIL)$
PCTINDUSED	$INODEUSED * 100.0 / f_files$
FSTYPE	getmntent API; struct mntent.mnt_type
PCTSPCAV	$100 - PCTSPCUSED$
MOUNTPTU	getmntent API; struct mntent.mnt_dir
PCTINDAVAL	$100 - PCTINDUSED$

LNXDSKIO attributes

The following table lists the mechanisms used to gather LNXDSKIO attributes.

Table 46. Mechanisms used to gather LNXDSKIO attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
TPS	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); reads + writes; $samples_("N" - "N - 1") / sample_interval$
BLKRDSSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); sectors read; $samples_("N" - "N - 1") / sample_interval$
BLKWRTNSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); sectors written; $samples_("N" - "N - 1") / sample_interval$
BLKSRD	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); total sectors read
BLKSWRTN	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); total sectors written

Table 46. Mechanisms used to gather LNXDSKIO attributes (continued)

Attribute name	Collection method
DEVMAJOR	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel)
DEVMINOR	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel)
DKNAME	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel)

LNxDU attributes

The following table lists the mechanisms used to gather LNxDU attributes.

Table 47. Mechanisms used to gather LNxDU attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
DSKNAME	getmntent API; struct mntent.mnt_fsname
SPCUSED	statfs API; struct statfs elements: $((f_blocks - f_bfree) * (f_bsize / 1024)) / 1024$
SPCAVAIL	statfs API; struct statfs elements: $(f_bavail * (f_bsize / 1024)) / 1024$
DURATE	Calculated from "N" and "N - 1" samples of SPCUSED
HWDURATE	Larger of "N" and "N - 1" samples of DURATE
HWTIME	Timestamp associated with the HWDURATE sample
DUMVAVG	Average of all DURATE values
DAYSDDK	$(SPCAVAIL * 1024 * 1024) / (DUMVAVG * 24)$
DAYSCURR	$(SPCAVAIL * 1024 * 1024 / (DURATE * 24))$
LWCURR	Smaller of "N" and "N - 1" samples of DAYSCURR
DAYSPEAK	$(SPCAVAIL * 1024 * 1024) / (HWDURATE * 24)$

LNxFILCMP attributes

The following table lists the mechanisms used to gather LNxFILCMP attributes.

Table 48. Mechanisms used to gather LNxFILCMP attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
FILENAME1	Passed to agent as parameter
FILENAME2	Passed to agent as parameter
COMPOPT	Passed to agent as parameter
COMPRESULT	Result from /usr/bin/cmp or /usr/bin/diff commands

LNXFILPAT attributes

The following table lists the mechanisms used to gather LNXFILPAT attributes.

Table 49. Mechanisms used to gather LNXFILPAT attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
PATHU	stat64, opendir and readdir APIs
FILEU	stat64, opendir and readdir APIs
SIZEMB	lstat64 API; struct stat64.st_size / (1024 * 1024) * 1000
OWNERU	lstat64 API; struct stat64.st_uid & getpwuid API - struct passed.pw_name
GROUPU	lstat64 API; struct stat64.st_gid & getgrgid API - struct group.gr_name
CHANGEDTM	lstat64 API; struct stat64.st_mtime
ACCESSEDTM	lstat64 API; struct stat64.st_atime
LINKS	lstat64 API; struct stat64.st_nlinks
ACCESS	lstat64 API; struct stat64.st_mode
TYPE	lstat64 API; struct stat64.st_mode
LINKNAMEU	readlink API
MODE	lstat64 API; struct stat64.st_mode
STATUSTM	lstat64 API; struct stat64.st_ctime
HASHALGO	Passed to agent as parameter
HASHSUM	CRC32, MD5 or SHA1 calculation - internal functions
FCCHANGED	Generated; true if HASHSUM has changed since last monitoring interval
SIZEMB64	lstat64 API; struct stat64.st_size / (1024 * 1024) * 1000

LNXFILPAT attributes

The following table lists the mechanisms used to gather LNXFILPAT attributes.

Table 50. Mechanisms used to gather LNXFILPAT attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
FILENAME	Passed to agent as parameter
MATCHPAT	Passed to agent as parameter
MATCHOPT	Passed to agent as parameter
MATCHCNT	Result from grep cmd

LNKGROUP attributes

The following table lists the mechanisms used to gather LNKGROUP attributes.

Table 51. Mechanisms used to gather LNKGROUP attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
GRPNAME	getgrouper API; struct group.gr_name
GRPID	getgrouper API; struct group.gr_gid
GRPDUP	Generated; true if duplicate GRPNAME or GRPID is detected
GRPID64	

LNKIOEXT attributes

The following table lists the mechanisms used to gather LNKIOEXT attributes.

Table 52. Mechanisms used to gather LNKIOEXT attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
DKNAME	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel)
RDRQMSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); reads merged; samples_("N" - "N - 1") / sample_interval
WRTRQMSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); writes merged; samples_("N" - "N - 1") / sample_interval
RDRQSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); reads; samples_("N" - "N - 1") / sample_interval
WRTREQSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); writes; samples_("N" - "N - 1") / sample_interval
RDSECTSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); sectors read; samples_("N" - "N - 1") / sample_interval
WRSECTSEC	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); sectors written; samples_("N" - "N - 1") / sample_interval
AVGRQSZ	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); samples_("N" - "N - 1"); (sectors_read + sectors_written) / (totals_reads + total_writes)
AVGRQQUSZ	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); I/O in progress; samples_("N" - "N - 1") / sample_interval
AVGWAITTM	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); samples_("N" - "N - 1"); (time_reading + time_writing) / (totals_reads + total_writes)
AVGSVCTM	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); samples_("N" - "N - 1"); time_spent_on_I/O / (totals_reads + total_writes)
IOUTIL	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); samples_("N" - "N - 1"); time_spent_on_I/O / monitoring_interval

Table 52. Mechanisms used to gather LNXIOEXT attributes (continued)

Attribute name	Collection method
IUTIL	Read from /proc/diskstats (2.6 kernel) or /proc/partitions (2.4 kernel); samples_("N" - "N - 1");IOUTIL / (total_reads / (totals_reads + total_writes))
OUTIL	IOUTIL - OUTIL
RDBYTESEC	RDSECTSEC converted to bytes
WRBYTESEC	WRSECTSEC converted to bytes
TOTBTSEC	WRBYTESEC + RDBYTESEC

LNXIPADDR attributes

The following table lists the mechanisms used to gather LNXIPADDR attributes.

Table 53. Mechanisms used to gather LNXIPADDR attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
INTFNAME	Read from /proc/net/dev
IPADDRESS	IPv4: socket, ioctl & inet_ntoa APIs. IPv6: read from /proc/net/if_inet6
DNSNAME	getaddrinfo and getnameinfo APIs
IPVERSION	Hardcoded based on IP type

LNXLOGIN attributes

The following table lists the mechanisms used to gather LNXLOGIN attributes.

Table 54. Mechanisms used to gather LNXLOGIN attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
USRNAME	getutent API; struct utmp.ut_user
USRPID	getutent API; struct utmp.ut_pid
LINE	getutent API; struct utmp.ut_line
LOGINTIME	getutent API; struct utmp.ut_tv.tv_sec
IDLETIME	stat API on /dev/ut_line to get last access time & subtract from current time
FROMHOST	getutent API; struct utmp.ut_host
USRNAMEU	getutent API; struct utmp.ut_user

LNXMACHIN attributes

The following table lists the mechanisms used to gather LNXMACHIN attributes.

Table 55. Mechanisms used to gather LNXMACHIN attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
HOSTNAME	gethostname API
BRAND	dmidecode program, where supported; read from /proc/sysinfo on zLinux; hardcoded on LinuxPPC
MODEL	dmidecode program, where supported; read from /proc/sysinfo on zLinux; read from /proc/device-tree/model on LinuxPPC
ONLNCPU	sysconf API
CONFPCPU	sysconf API
BIOSVER	dmidecode program, where supported; hardcoded otherwise
BIOSREL	dmidecode program, where supported; hardcoded otherwise
MACSERIAL	dmidecode program, where supported; read from /proc/sysinfo on zLinux; read from /proc/device-tree/system-id on LinuxPPC
UUID	dmidecode program, where supported; hardcoded otherwise
VMID	read from /proc/sysinfo on zLinux; read from /proc/ppc64/lparcfg on LinuxPPC; empty otherwise

LNXNET attributes

The following table lists the mechanisms used to gather LNXNET attributes.

Table 56. Mechanisms used to gather LNXNET attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
FNAME	Read from /proc/net/dev
FIPADDR	socket, ioctl & inet_ntoa APIs
FSTATUS	socket & ioctl APIs
FMTU	socket & ioctl APIs
FIKBYTES	Read from /proc/net/dev & divided by 1024
RECBPS	Read from /proc/net/dev; samples_("N" - "N - 1") / sample_interval
FOKBYTES	Read from /proc/net/dev & divided by 1024
TRANSBPS	Read from /proc/net/dev; samples_("N" - "N - 1") / sample_interval
FIFRAMES	Read from /proc/net/dev
RPACKPS	Read from /proc/net/dev; samples_("N" - "N - 1") / sample_interval
FOFRAMES	Read from /proc/net/dev
TPACKPS	Read from /proc/net/dev; samples_("N" - "N - 1") / sample_interval

Table 56. Mechanisms used to gather LNXNET attributes (continued)

Attribute name	Collection method
FIERRORS	Read from /proc/net/dev
FOERRORS	Read from /proc/net/dev
FCOLLSNS	Read from /proc/net/dev
FCOLLSNRT	Read from /proc/net/dev; samples_("N" - "N - 1") * 60 / sample_interval
FCOLLSPCT	Read from /proc/net/dev; for this sample period: (collisions / (frames sent + frames rcved)) * 100
FIERRORT	Read from /proc/net/dev; samples_("N" - "N - 1") * 60 / sample_interval
FOERRORT	Read from /proc/net/dev; samples_("N" - "N - 1") * 60 / sample_interval
FIOERPCT	Read from /proc/net/dev; for this sample period: (input_errors + output_errors) / (frames_sent + frames_rcved) * 100
FIDROP	Read from /proc/net/dev
FODROP	Read from /proc/net/dev
FIFOINOV	Read from /proc/net/dev
FIFOUTOVR	Read from /proc/net/dev
FIPKTFRAM	Read from /proc/net/dev
FCARRIER	Read from /proc/net/dev
FIERPCT	FIOERPCT * (FIERRORT / (FIERRORT + FOERRORT))
FOERPCT	FIOERPCT - FIERPCT
DEVTYPE	socket & ioctl APIs
MACADDRESS	socket & ioctl APIs

LNXNFS attributes

The following table lists the mechanisms used to gather LNXNFS attributes.

Table 57. Mechanisms used to gather LNXNFS attributes

Attribute name	Collection method
LOCORIG	Generated client/server indicator
NFSVER	Generated version indicator
NFSNULL	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
NULLPCT	NFSNULL * 100 / NFSTOT
NFSGETATT	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
GETATTPCT	NFSGETADD * 100 / NFSTOT
NFSSETATT	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
SETATTPCT	NFSSETATT * 100 / NFSTOT
NFSROOT	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
NFSROOTPCT	NFSROOT * 100 / NFSTOT

Table 57. Mechanisms used to gather LNXNFS attributes (continued)

Attribute name	Collection method
NFSLOOKUP	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
LOOKUPPCT	NFSLOOKUP * 100 / NFSTOT
NFSRDLINK	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
RDLINKPCT	NFSRDLINK * 100 / NFSTOT
NFSREAD	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
READPCT	NFSREAD * 100 / NFSTOT
NFSWRCACH	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
WRCACHPCT	NFSWRCACH * 100 / NFSTOT
NFSWRITES	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
NFSWRTPCT	NFSWRITES * 100 / NFSTOT
NFSCREATES	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
CREATESPCT	NFSCREATES * 100 / NFSTOT
NFSREMOVE	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
REMOVEPCT	NFSREMOVE * 100 / NFSTOT
NFSRENAME	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
RENMPCT	NFSRENAME * 100 / NFSTOT
NFSLINK	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
LINKPCT	NFSLINK * 100 / NFSTOT
NFSSYMLNK	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
SYMLNKPCT	NFSSYMLNK * 100 / NFSTOT
NFSMKDIR	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
MKDIRPCT	NFSMKDIR * 100 / NFSTOT
NSRMDIR	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
RMDIRPCT	NFSRMDIR * 100 / NFSTOT
NFSRDDIR	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
RDDIRPCT	NFSRDDIR * 100 / NFSTOT
NFSFSSTAT	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
FSSTATPCT	NFSFSSTAT * 100 / NFSTOT
NFSACCESS	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)

Table 57. Mechanisms used to gather LNXNFS attributes (continued)

Attribute name	Collection method
ACCSSPCT	NFSACCESS * 100 / NFSTOT
NFSMKNOD	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
MKNODPCT	NFSMKNOD * 100 / NFSTOT
RDDIRPLUS	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
RDIRPLSPCT	RDDIRPLUS * 100 / NFSTOT
NFSFSINFO	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
FSINFOPCT	NFSFSINFO * 100 / NFSTOT
NFSPHCONF	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
PTHCONFPCT	NFSPHCONF * 100 / NFSTOT
NFSCOMMIT	Read from /proc/net/rpc/nfs (client) or /proc/net/rpc/nfsd (server)
COMMITPCT	NFSCOMMIT * 100 / NFSTOT
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
NFSTOT	"NFSNULL + NFSGETATT + NFSSETATT + NFSROOT + NFSLOOKUP + NFSRDLINK + NFSREAD + NFSWRKACH + NFSWRITES + NFScreates + NFSREMOVE + NFSRENAME + NFSLINK + NFSYMLNK + NFSMKDIR + NFSRMDIR + NFSRDIR + NFSSTAT"

LNXOSCON attributes

The following table lists the mechanisms used to gather LNXOSCON attributes.

Table 58. Mechanisms used to gather LNXOSCON attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
OSNAME	Read from /proc/version
OSVER	Read from /proc/version
GCCVER	Read from /proc/version
OSVEND	Read from /proc/version

LNXPING attributes

The following table lists the mechanisms used to gather LNXPING attributes.

Table 59. Mechanisms used to gather LNXPING attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time

Table 59. Mechanisms used to gather LNXPING attributes (continued)

Attribute name	Collection method
TGTSERV	Passed to agent as situation parameter or from KLZ_PINGHOSTLIST file
SERVUP	Result from /bin/ping command
HOSTRESP	Result from /bin/ping command

LNXPROC attributes

The following table lists the mechanisms used to gather LNXPROC attributes.

Table 60. Mechanisms used to gather LNXPROC attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
PID	Read from /proc; PID is the subdirectory name
PPID	Read from /proc/PID/stat
CMD	Read from /proc/PID/stat
STATE	Read from /proc/PID/stat
PSYSCPU	Read from /proc/PID/stat; converted from jiffies
PUSRCPU	Read from /proc/PID/stat; converted from jiffies
TSYSCPU	Read from /proc/PID/stat; converted from jiffies
TUSRCPU	Read from /proc/PID/stat
INTPRI	Read from /proc/PID/stat
NICE	Read from /proc/PID/statm
SIZE	Read from /proc/PID/statm
RSS	Read from /proc/PID/statm
SHAREMEM	Read from /proc/PID/statm
TRS	Read from /proc/PID/statm
LRS	Read from /proc/PID/statm
DRS	Read from /proc/PID/statm
DIRTPG	Read from /proc/PID/statm
VMSIZE	Read from /proc/PID/status
VMLOCK	Read from /proc/PID/status
VMDATA	Read from /proc/PID/status
VMSTACK	Read from /proc/PID/status
VMEXESZ	Read from /proc/PID/status
VMLIBSZ	Read from /proc/PID/status
CMINFLT	Read from /proc/PID/stat
CMAJFLT	Read from /proc/PID/stat
CMDLINE	Read from /proc/PID/cmdline
CMDLINEU	Read from /proc/PID/cmdline
CPUAFF	Read from /proc/PID/stat

Table 60. Mechanisms used to gather LNXPROC attributes (continued)

Attribute name	Collection method
USRSYSCPU	(TUSRCPU / TSYSCPU) * 100
CMDU	Read from /proc/PID/stat
TBUSYCPU	TSYSCPU + TUSRCPU
BUSYCPU	PSYSCPU + PUSRCPU
VMSIZEMB	Read from /proc/PID/status; converted to MB
VMLOCKMB	Read from /proc/PID/status; converted to MB
VMDATAMB	Read from /proc/PID/status; converted to MB
VMSTACKMB	Read from /proc/PID/status; converted to MB
VMEXESZMB	Read from /proc/PID/status; converted to MB
VMLIBSZMB	Read from /proc/PID/status; converted to MB
PROCTHRD	Read from /proc/PID/status
SESSIONID	Read from /proc/PID/stat
PSYSNORM	Read from /proc/PID/stat; converted from jiffies
PUSRNORM	Process user-mode time read from /prod/PID/stat; Nbr of CPUs read from sysconf API; (current_user_mode - previous_user_mode) / (elapsed_time * nbr_of_cpus)
PBUSYNORM	Process kernel-mode time read from /prod/PID/stat; Nbr of CPUs read from sysconf API; (current_kernel_mode - previous_kernel_mode) / (elapsed_time * nbr_of_cpus)
PROCCOUNT	Generated; count of processes with same CMDLINE

LNXPUSE attributes

The following table lists the mechanisms used to gather LNXPUSE attributes.

Table 61. Mechanisms used to gather LNXPUSE attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
PID	Read from /proc; PID is the subdirectory name
RUSER	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
EUSER	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
SUSER	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
FSUSER	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
RGRP	Read from /proc/PID/status; converted to string using the getgrgid API; struct group.gr_name
EGRP	Read from /proc/PID/status; converted to string using the getgrgid API; struct group.gr_name
SGRP	Read from /proc/PID/status; converted to string using the getgrgid API; struct group.gr_name

Table 61. Mechanisms used to gather LNXUSER attributes (continued)

Attribute name	Collection method
FSGRP	Read from /proc/PID/status; converted to string using the getgrgid API; struct group.gr_name
RUSERID	Read from /proc/PID/status
EUSERID	Read from /proc/PID/status
SUSERID	Read from /proc/PID/status
FSUSRID	Read from /proc/PID/status
RGRPID	Read from /proc/PID/status
EFFGRPID	Read from /proc/PID/status
SGRPID	Read from /proc/PID/status
FSGRPID	Read from /proc/PID/status
RUSERU	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
EUSERU	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
SUSERU	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
FSUSERU	Read from /proc/PID/status; converted to string using the getpwuid API; struct passwd.pw_name
RGRPU	Read from /proc/PID/status; converted to string using the getgrgid API; struct group.gr_name
EGRPU	Read from /proc/PID/status; converted to string using the getgrgid API; struct group.gr_name
SGRPU	Read from /proc/PID/status; converted to string using the getgrgid API; struct group.gr_name
FSGRPU	Read from /proc/PID/status; converted to string using the getgrgid API; struct group.gr_name
SESSIONID	Read from /proc/PID/stat
PPID	Read from /proc/PID/stat
STATE	Read from /proc/PID/stat
CMDLINEU	Read from /proc/PID/cmdline
CMDU	Read from /proc/PID/stat
VMSIZEMB	Read from /proc/PID/status; converted to MB
TTY	Read from /proc/PID/stat; converted to string by internal method

LNXRPC attributes

The following table lists the mechanisms used to gather LNXRPC attributes.

Table 62. Mechanisms used to gather LNXRPC attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
RSCALLS	Read from /proc/net/rpc/nfsd
RSBADCALL	Read from /proc/net/rpc/nfsd

Table 62. Mechanisms used to gather LNXRPC attributes (continued)

Attribute name	Collection method
RSBADAUTH	Read from /proc/net/rpc/nfsd
RSBADCLT	Read from /proc/net/rpc/nfsd
RSXDRCALL	Read from /proc/net/rpc/nfsd
RCCALLS	Read from /proc/net/rpc/nfs
RCRETRAN	Read from /proc/net/rpc/nfs
RCAREF	Read from /proc/net/rpc/nfs

LNXSOCKD attributes

The following table lists the mechanisms used to gather LNXSOCKD attributes.

Table 63. Mechanisms used to gather LNXSOCKD attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
SCKPROTO	Generated TCP/UDP indicator
RECVQ	Read from /proc/net/tcp or /proc/net/udp
SENDQ	Read from /proc/net/tcp or /proc/net/udp
LOCLADDR	Read from /proc/net/tcp or /proc/net/udp
LOCLPORT	Read from /proc/net/tcp or /proc/net/udp
LOCLSVC	Read from /proc/net/tcp or /proc/net/udp & getservbyport API; struct servent.s_name
FORNADDR	Read from /proc/net/tcp or /proc/net/udp
STATE	Read from /proc/net/tcp or /proc/net/udp
SOCKUID	Read from /proc/net/tcp or /proc/net/udp
SCKINOD	Read from /proc/net/tcp or /proc/net/udp
REMOTPORT	Read from /proc/net/tcp or /proc/net/udp
RUSERU	Read from /proc/net/tcp or /proc/net/udp & getpwuid API; struct passwd.pw_name

LNXSOCKS attributes

The following table lists the mechanisms used to gather LNXSOCKS attributes.

Table 64. Mechanisms used to gather LNXSOCKS attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
SCKPROTO	Read from /proc/net/sockstat
SCKINUSE	Read from /proc/net/sockstat
SCKHWUSED	Read from /proc/net/sockstat

LNXPRT attributes

The following table lists the mechanisms used to gather LNXPRT attributes.

Table 65. Mechanisms used to gather LNXPRT attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
MOVSWPTOT	Read from /proc/meminfo; (last MOVSWAPTOT + SwapTotal) / 2
SWAPUSED	Read from /proc/meminfo; (last SWAPUSED + (SwapTotal - SwapFree)) / 2
SWPRATE	Read from /proc/meminfo; (last SWAPRATE + ((SwapTotal - SwapFree) - previous_SWAPUSED)) / 2
SWAPDAYS	Read from /proc/meminfo; SwapTotal / (24 * SWAPRATE)
PKSWPUSD	Read from /proc/meminfo; larger of last two (SwapTotal - SwapFree)
MINSWPDAYS	Read from /proc/meminfo; smaller of last two SWAPDAYS
LOWMEM	Read from /proc/meminfo; LowFree

LNXSYS attributes

The following table lists the mechanisms used to gather LNXSYS attributes.

Table 66. Mechanisms used to gather LNXSYS attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
CSWSEC	Read from /proc/stat; samples_("N" - "N - 1") / sample_interval
RATECSW	Read from /proc/stat; ((current_CSWSEC - previous_CSWSEC) / previous_CSWSEC) * 100
PROCSEC	Read from /proc/stat; samples_("N" - "N - 1") / sample_interval
RATEPROC	Read from /proc/stat; ((current_PROCSEC - previous_PROCSEC) / previous_PROCSEC) * 100
CURUSRS	getutent API; count of entries in utmp database
LOAD1MIN	Read from /proc/loadavg * 100
LOAD5MIN	Read from /proc/loadavg * 100
LOAD15MIN	Read from /proc/loadavg * 100
SYSUPTIME	Read from /proc/uptime
PGPGIN	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel)
PGPGINPS	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel); samples_("N" - "N - 1") / sample_interval * 100
PGPGOUT	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel)
PGPGOUTPS	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel); samples_("N" - "N - 1") / sample_interval * 100
PGSWAPIN	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel)
SWAPINPS	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel); samples_("N" - "N - 1") / sample_interval * 100

Table 66. Mechanisms used to gather LNXSYS attributes (continued)

Attribute name	Collection method
PGSWAPOUT	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel)
SWAPOUTPS	Read from /proc/vmstat (2.6 kernel) or /proc/stat (2.4 kernel); samples_("N" - "N - 1") / sample_interval * 100
PGFLTPTS	Read from /proc/vmstat (2.6 kernel) samples_("N" - "N - 1") / sample_interval * 100; N/A on 2.4 kernel
MAJFLTPTS	Read from /proc/vmstat (2.6 kernel) samples_("N" - "N - 1") / sample_interval * 100; N/A on 2.4 kernel
TOTPROCS	Count process subdirs in /proc
ZOMBCNT	Count process subdirs in /proc in zombie state

LNXVM attributes

The following table lists the mechanisms used to gather LNXVM attributes.

Table 67. Mechanisms used to gather LNXVM attributes

Attribute name	Collection method
ORIGINNODE	Short host name + ":LZ"
TIMESTAMP	Current time
SWAPTOT	Read from /proc/meminfo; (SwapTotal / 1024) * 100
SWAPUSED	Read from /proc/meminfo; ((SwapTotal - SwapFree) / 1024) * 100
SWAPFREE	Read from /proc/meminfo; (SwapFree / 1024) * 100
MEMTOT	Read from /proc/meminfo; (MemTotal / 1024) * 100
MEMUSED	Read from /proc/meminfo; ((MemTotal - MemFree) / 1024) * 100
MEMFREE	Read from /proc/meminfo; (MemFree / 1024) * 100
MEMSHARED	Read from /proc/meminfo; (MemShared / 1024) * 100
MEMBUFF	Read from /proc/meminfo; (Buffers / 1024) * 100
MEMCACHE	Read from /proc/meminfo; (Cache / 1024) * 100
VSTOT	MEMTOT + SWAPTOT
USEDVVS	SWAPUSED + MEMUSED
AVAILVS	VSTOT - USEDVVS
AVALVSPCT	(AVAILVS / VSTOT) * 100
USEDVSPCT	100 - USEDVSPPCT
USEDRLPCT	Read from /proc/meminfo; ((MemTotal - MemFree) / MemTotal) * 100
AVALRLPCT	100 - USEDVSPPCT
USEDVSPPCT	Read from /proc/meminfo; ((SwapTotal - SwapFree) / SwapTotal) * 100
AVALSWPPCT	100 - USEDVSPPCT

Appendix D. Discovery Library Adapter for the monitoring agent

This chapter contains information about the Discovery Library Adapter (DLA) for the Monitoring Agent for Linux.

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 Linux resources. For all Linux 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 Linux components.

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

Linux 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 Linux.

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

Linux class

The following information describes the Linux class.

CDM class name

sys.linux.Linux or sys.zOS.ZLinux

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: OSVER/LNXOSCON
Description: OS Version
- CDM attribute: Name
Agent attribute: none
Description: "Linux"
- CDM attribute: FQDN
Agent attribute: DNSNAME/LNXIPADDR
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/LNXIPADDR and MACADDRESS/LNXNET
Description: Lowest IP Address (MAC Address)
- CDM attribute: PrimaryMACAddress
Agent attribute: MACADDRESS/LNXNET
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/LNXIPADDR
Description: Fully Qualified Domain Name
- CDM attribute: SystemBoardUUID
Agent attribute: UUID/LNXMACHIN
Description: System Board UUID
- CDM attribute: SerialNumber
Agent attribute: MACSERIAL/LNXMACHIN
Description: Serial Number
- CDM attribute: Model
Agent attribute: MODEL/LNXMACHIN
Description: Model
- CDM attribute: Manufacturer

Agent attribute: BRAND/LNXMACHIN

Description: Manufacturer

- 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
IPADDRESS/LNXIPADDR
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: DNSNAME/LNXIPADDR
Description: Fully Qualified Domain Name

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 (LZ)
- CDM attribute: Affinity
Description: OS Agent affinity
- CDM attribute: Label
Description: Managed System Name "- Linux OS"

Appendix E. OSLC resources for the Linux OS agent

Tivoli Monitoring includes the Open Services for Lifecycle Collaboration Performance Monitoring (OSLC-PM) service provider. The service provider registers monitoring resources with the Registry Services. Registry Services is a Jazz for Service Management integration service that provides a shared data repository for products in an integrated service management environment. Products that discover and manage shared IT resources can register these IT resources and the services they offer with Registry Services. Other products such as Tivoli Business Service Manager can use data by querying Registry Services for the managed resources or the associated service providers of interest.

The OSLC Performance Monitoring service provider is included with the Tivoli Enterprise Monitoring Automation Server. The automation server is installed on the same computer system as the Hub monitoring server and is configured with the connection information for Registry Services. The Performance Monitoring service provider registers resource types that are defined by the OSLC Common Resource Type Vocabulary (CRTV). Monitoring agents provide a template that maps their monitoring data to CRTV resources. The template is installed with the agent's monitoring server application support in these directories on the computer system where the automation server and Hub monitoring server are installed:

- For UNIX or Linux: CANDLEHOME/tables/cicatrsq/OSLC/xml
- For Windows or IBM i: CANDLEHOME/CMS/OSLC/xml

The Linux OS agent's OSLC support is dependent on the type of Linux system:

- On Intel based systems, the Linux OS agent must be at version 6.2.2 or later.
- On PowerPC based systems or Linux on System z, the Linux OS agent must be at version 6.3 or later.

Note: The Linux OS agent's application support on the Hub monitoring server must be at version 6.3 or later.

The Performance Monitoring service provider registers ComputerSystem and IPAddress resources for OS agents with Registry Services when an agent comes online for the first time. The service provider also periodically checks for resource changes such as new IP addresses and notifies Registry Services of any updates.

The Performance Monitoring service provider also provides an OSLC RESTful API for retrieving linked data about ComputerSystem and IPAddress resources. The service provider accepts HTTP GET requests for the RDF/XML, compact XML, and HTML content types. However, it returns an HTTP 406 status code when it receives a compact XML request for IPAddress resources since UI preview (HTML content) is not supported for that resource type. Also, the HTML content is not translated so is always displayed in English.

When RDF/XML and HTML content is requested for a ComputerSystem resource, the service provider returns the properties that were registered with Registry Services and metric properties that are defined by the OSLC Performance Monitoring working group and the private ITM namespace vocabulary.

When RDF/XML content is requested for an IPAddress resource, only the registered properties are returned since metrics are not defined for this resource type.

The following sections list the CRTV resources and properties that are registered for OS agents and also the metric properties that the service provider returns in RDF/XML responses or that are used to return HTML content.

For more information on Common Resource Type Vocabulary resources and Performance Monitoring metrics, see:

- <http://open-services.net/wiki/reconciliation/>
- <http://open-services.net/wiki/performance-monitoring/>

ComputerSystem resources

The following information describes the ComputerSystem resource information that is provided by the Linux OS monitoring agent. The corresponding attributes, from which the properties are derived, are listed within parentheses.

Resource type

<http://open-services.net/ns/crtv#ComputerSystem>

Record type

<http://jazz.net/ns/ism/perfmon/itm#klz>

Relationships

None

Vocabulary namespaces

crtv=<http://open-services.net/ns/crtv#>

ems=<http://open-services.net/ns/ems#>

itm=<http://jazz.net/ns/ism/perfmon/itm#>

oslc=<http://open-services.net/ns/core#>

pm=<http://open-services.net/ns/perfmon#>

Registration record properties

crtv:fqdn (LNXMACHIN.HOSTNAME)

crtv:model (LNXMACHIN.MODEL)

crtv:manufacturer (LNXMACHIN.BRAND)

crtv:serialNumber (LNXMACHIN.MACSERIAL)

crtv:shortHostname (LNXMACHIN.HOSTNAME)

crtv:systemBoardUUID (LNXMACHIN.UUID)

crtv:vmid (LNXMACHIN.VMID)

itm:internalID (This property is not intended for use by OSLC client applications.)

itm:managedSystemName (LNXMACHIN.NODE)

itm:osType (INODESTS.HOSTINFO)

oslc:domain (<http://open-services.net/ns/perfmon#>)

oslc:serviceProvider (URL of the Performance Monitoring service provider record in Registry Services)

Note: The crtv:systemBoardUUID (LNXMACHIN.UUID) property is not registered for Linux systems on PowerPC or Linux on System z.

Note: The crtv:fqdn crt:vmid (LNXMACHIN.VMID) property is only registered for Linux systems on PowerPC or Linux on System z.

Metric properties

itm:disksByPercentageSpaceUsed (KLZDISK.MOUNTPT
KLZDISK.DSKUSEDPCT)
itm:monitoringAgentsByStatus (KLZPASSTAT.PASAGTNAME
KLZPASSTAT.INSTNAME KLZPASSTAT.STATUS)
itm:nativeIdentity (INODESTS.ORIGINNODE INODESTS.AFFINITIES)
itm:topProcessesforCPUUtil (KLZPROC.CMD KLZPROC.PID
KLZPROC.CPUPERCENT)
itm:topProcessesForRealMemUtil (KLZPROC.CMD KLZPROC.PID
KLZPROC.RSS)
itm:topProcessesforVirtMemUtil (KLZPROC.CMD KLZPROC.PID
KLZPROC.VMSIZE)

Note: Agent Management Services provides the list of installed agents and their respective status.

IPAddress resources

The following information describes the IPAddress resource information that is provided by the Linux OS monitoring agent.

Resource type

<http://open-services.net/ns/crtv#IPAddress>

Record type

<http://jazz.net/ns/ism/perfmon/itm#klz>

Relationships

dependsOn where the target is a resource type of <http://open-services.net/ns/crtv#ComputerSystem>

Vocabulary namespaces

oslc=<http://open-services.net/ns/core#>
crtv=<http://open-services.net/ns/crtv#>
itm=<http://jazz.net/ns/ism/perfmon/itm#>

Registration record properties

crtv:address (LNXIPADDR.IPADDRESS)
crtv:contextAddressSpace (<http://open-services.net/ns/crtv#NULL>)
itm:internalID (This property is not intended for use by OSLC client applications.)
itm:managedSystemName (LNXMACHIN.NODE)
oslc:domain (<http://open-services.net/ns/perfmon#>)
oslc:serviceProvider (URL of the Performance Monitoring service provider record in Registry Services)

Note: All IP addresses are registered with the public IP address context value.

Metric properties

None.

Sample RDF/XML response for a ComputerSystem resource

```
<rdf:RDF
  xml:base="http://10.10.45.34:10001/kas_srv/provider?type=ComputerSystem&
search=1352404171062505"
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:oslc="http://open-services.net/ns/core#"
  xmlns:dcterms="http://purl.org/dc/terms/"
```

```

xmlns:pm="http://open-services.net/ns/perfmon#"
xmlns:crtv="http://open-services.net/ns/crtv#"
xmlns:itm="http://jazz.net/ns/ism/perfmon/itm#"
xmlns:ems="http://open-services.net/ns/ems#"
<rdf:Description rdf:about="http://10.10.45.34:10001/kas_srv/provider?
type=ComputerSystem&search=1352404171062505">
  <crtv:model>911051A</crtv:model>
  <crtv:manufacturer>IBM</crtv:manufacturer>
  <crtv:shortHostname>nc117244</crtv:shortHostname>
  <itm:osType rdf:resource="http://jazz.net/ns/ism/perfmon/itm/osType#Linux"/>
  <crtv:fgdn>nc117244.rome1ab.it.ibm.com</crtv:fgdn>
  <crtv:vmid>7</crtv:vmid>
  <crtv:serialNumber>0651D12</crtv:serialNumber>
  <rdf:type rdf:resource="http://open-services.net/ns/perfmon#PerformanceMonitoringRecord"/>
  <dcterms:isPartOf rdf:resource="http://10.10.45.34:10001/kas_srv/provider?
type=ComputerSystem&search=1352404171062505"/>
  <dcterms:date>11/13/2012T18:10:39Z</dcterms:date>
  <itm:nativeIdentity>
  <rdf:Description>
    <itm:managedSystemName>nc117244:LZ</itm:managedSystemName>
    <itm:affinity>%IBM.STATIC134      000000000000Jyw0a7</itm:affinity>
  </rdf:Description>
  </itm:nativeIdentity>
  <itm:managedSystemName>nc117244:LZ</itm:managedSystemName>
  <oslc:domain rdf:resource="http://open-services.net/ns/perfmon#" />
  <oslc:serviceProvider rdf:resource="http://nc112031:16310/oslc/providers/1351248331634"/>
  <itm:topProcessesForCPUUtil rdf:resource="#topCPU-seq"/>
  <itm:topProcessesForVirtMemUtil rdf:resource="#topVirtMem-seq"/>
  <itm:topProcessesForRealMemUtil rdf:resource="#topRealMem-seq"/>
  <itm:disksByPercentageSpaceUsed rdf:resource="#diskSpaceUsed-seq"/>
  <itm:monitoringAgentsByStatus rdf:resource="#monitoringAgentsStatus-seq"/>
</rdf:Description>
<rdf:Description rdf:about="#topCPU-seq">
  <rdf:type rdf:resource="http://www.w3.org/1999/02/22-rdf-syntax-ns#Seq"/>
  <rdf:_1 rdf:resource="#klzagent-27517"/>
  <rdf:_2 rdf:resource="#klzagent-20399"/>
</rdf:Description>

<rdf:Description rdf:about="#topVirtMem-seq">
  <rdf:type rdf:resource="http://www.w3.org/1999/02/22-rdf-syntax-ns#Seq"/>
  <rdf:_1 rdf:resource="#klzagent-24848"/>
  <rdf:_2 rdf:resource="#klzagent-27517"/>
  <rdf:_3 rdf:resource="#klzagent-3430"/>
  <rdf:_4 rdf:resource="#klzagent-7293"/>
  <rdf:_5 rdf:resource="#klzagent-20399"/>
</rdf:Description>

<rdf:Description rdf:about="#topRealMem-seq">
  <rdf:type rdf:resource="http://www.w3.org/1999/02/22-rdf-syntax-ns#Seq"/>
  <rdf:_1 rdf:resource="#klzagent-24848"/>
  <rdf:_2 rdf:resource="#klzagent-27517"/>
  <rdf:_3 rdf:resource="#klzagent-7293"/>
  <rdf:_4 rdf:resource="#klzagent-3430"/>
  <rdf:_5 rdf:resource="#klzagent-20399"/>
</rdf:Description>

<rdf:Description rdf:about="#diskSpaceUsed-seq">
  <rdf:type rdf:resource="http://www.w3.org/1999/02/22-rdf-syntax-ns#Seq"/>
  <rdf:_1 rdf:resource="#"/>
  <rdf:_2 rdf:resource="#/dev"/>
</rdf:Description>

<rdf:Description rdf:about="#monitoringAgentsStatus-seq">
  <rdf:type rdf:resource="http://www.w3.org/1999/02/22-rdf-syntax-ns#Seq"/>
  <rdf:_1 rdf:resource="#Monitoring%20Agent%20for%20Linux%200S-None"/>
  <rdf:_2 rdf:resource="#Proxy%20Agent%20Services%20Watchdog"/>
</rdf:Description>

<rdf:Description rdf:about="#klzagent-27517">
  <rdf:type rdf:resource="http://open-services.net/ns/crtv#Process"/>
  <rdf:type rdf:resource="http://open-services.net/ns/perfmon#PerformanceMonitoringRecord"/>
  <dcterms:isPartOf rdf:resource="http://10.10.45.34:10001/kas_srv/provider?
type=ComputerSystem&search=1352404171062505"/>
  <dcterms:title>klzagent</dcterms:title>
  <crtv:processId>27517</crtv:processId>
  <ems:observes rdf:resource="#klzagent-27517-cpuUtil"/>
  <ems:observes rdf:resource="#klzagent-27517-virtMem"/>
  <ems:observes rdf:resource="#klzagent-27517-realMem"/>
</rdf:Description>

<rdf:Description rdf:about="#klzagent-20399">
  <rdf:type rdf:resource="http://open-services.net/ns/crtv#Process"/>
  <rdf:type rdf:resource="http://open-services.net/ns/perfmon#PerformanceMonitoringRecord"/>
  <dcterms:isPartOf rdf:resource="http://10.10.45.34:10001/kas_srv/provider?
type=ComputerSystem&search=1352404171062505"/>
  <dcterms:title>klzagent</dcterms:title>
  <crtv:processId>20399</crtv:processId>
  <ems:observes rdf:resource="#klzagent-20399-cpuUtil"/>
  <ems:observes rdf:resource="#klzagent-20399-virtMem"/>
  <ems:observes rdf:resource="#klzagent-20399-realMem"/>
</rdf:Description>

<rdf:Description rdf:about="#klzagent-24848">
  <rdf:type rdf:resource="http://open-services.net/ns/crtv#Process"/>
  <rdf:type rdf:resource="http://open-services.net/ns/perfmon#PerformanceMonitoringRecord"/>
  <dcterms:isPartOf rdf:resource="http://10.10.45.34:10001/kas_srv/provider?
type=ComputerSystem&search=1352404171062505"/>

```

```

    <dcterms:title>klzagent</dcterms:title>
    <crtv:processId>24848</crtv:processId>
    <ems:observes rdf:resource="#klzagent-24848-virtMem"/>
    <ems:observes rdf:resource="#klzagent-24848-realMem"/>
  </rdf:Description>
  <rdf:Description rdf:about="#klzagent-3430">
    <rdf:type rdf:resource="http://open-services.net/ns/crtv#Process"/>
    <rdf:type rdf:resource="http://open-services.net/ns/perfmon#PerformanceMonitoringRecord"/>
    <dcterms:isPartOf rdf:resource="http://10.10.45.34:10001/kas_srv/provider?
type=ComputerSystem&search=1352404171062505"/>
    <dcterms:title>klzagent</dcterms:title>
    <crtv:processId>3430</crtv:processId>
    <ems:observes rdf:resource="#klzagent-3430-virtMem"/>
    <ems:observes rdf:resource="#klzagent-3430-realMem"/>
  </rdf:Description>
  <rdf:Description rdf:about="#klzagent-7293">
    <rdf:type rdf:resource="http://open-services.net/ns/crtv#Process"/>
    <rdf:type rdf:resource="http://open-services.net/ns/perfmon#PerformanceMonitoringRecord"/>
    <dcterms:isPartOf rdf:resource="http://10.10.45.34:10001/kas_srv/provider?
type=ComputerSystem&search=1352404171062505"/>
    <dcterms:title>klzagent</dcterms:title>
    <crtv:processId>7293</crtv:processId>
    <ems:observes rdf:resource="#klzagent-7293-virtMem"/>
    <ems:observes rdf:resource="#klzagent-7293-realMem"/>
  </rdf:Description>
  <rdf:Description rdf:about="#"/>
    <rdf:type rdf:resource="http://open-services.net/ns/crtv#StorageVolume"/>
    <rdf:type rdf:resource="http://open-services.net/ns/perfmon#PerformanceMonitoringRecord"/>
    <dcterms:isPartOf rdf:resource="http://10.10.45.34:10001/kas_srv/provider?
type=ComputerSystem&search=1352404171062505"/>
    <dcterms:title></dcterms:title>
    <ems:observes rdf:resource="#diskUsed-"/>
  </rdf:Description>
  <rdf:Description rdf:about="#/dev">
    <rdf:type rdf:resource="http://open-services.net/ns/crtv#StorageVolume"/>
    <rdf:type rdf:resource="http://open-services.net/ns/perfmon#PerformanceMonitoringRecord"/>
    <dcterms:isPartOf rdf:resource="http://10.10.45.34:10001/kas_srv/provider?
type=ComputerSystem&search=1352404171062505"/>
    <dcterms:title>/dev</dcterms:title>
    <ems:observes rdf:resource="#diskUsed-/dev"/>
  </rdf:Description>
  <rdf:Description rdf:about="#Monitoring%20Agent%20for%20Linux%20S-None">
    <rdf:type rdf:resource="http://xmlns.com/foaf/0.1/Agent"/>
    <rdf:type rdf:resource="http://open-services.net/ns/perfmon#PerformanceMonitoringRecord"/>
    <dcterms:isPartOf rdf:resource="http://10.10.45.34:10001/kas_srv/provider?
type=ComputerSystem&search=1352404171062505"/>
    <dcterms:title>Monitoring Agent for Linux OS</dcterms:title>
    <pm:availabilityStatus rdf:resource="http://open-services.net/ns/perfmon#Running"/>
    <itm:instanceName>None</itm:instanceName>
  </rdf:Description>
  <rdf:Description rdf:about="#Proxy%20Agent%20Services%20Watchdog->
    <rdf:type rdf:resource="http://xmlns.com/foaf/0.1/Agent"/>
    <rdf:type rdf:resource="http://open-services.net/ns/perfmon#PerformanceMonitoringRecord"/>
    <dcterms:isPartOf rdf:resource="http://10.10.45.34:10001/kas_srv/provider?
type=ComputerSystem&search=1352404171062505"/>
    <dcterms:title>Proxy Agent Services Watchdog</dcterms:title>
    <pm:availabilityStatus rdf:resource="http://open-services.net/ns/perfmon#Running"/>
    <itm:instanceName></itm:instanceName>
  </rdf:Description>
  <rdf:Description rdf:about="#klzagent-27517-cpuUtil">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>CPU Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">0.05</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#CpuUsed"/>
  </rdf:Description>
  <rdf:Description rdf:about="#klzagent-20399-cpuUtil">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>CPU Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">0.02</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#CpuUsed"/>
  </rdf:Description>
  <rdf:Description rdf:about="#klzagent-24848-virtMem">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Virtual Memory Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">20.69</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#VirtualMemoryUsed"/>
  </rdf:Description>
  <rdf:Description rdf:about="#klzagent-27517-virtMem">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Virtual Memory Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">20.58</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#VirtualMemoryUsed"/>
  </rdf:Description>
  <rdf:Description rdf:about="#klzagent-3430-virtMem">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Virtual Memory Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">20.14</ems:numericValue>

```

```

        <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
        <ems:metric rdf:resource="http://open-services.net/ns/perfmon#VirtualMemoryUsed"/>
    </rdf:Description>
<rdf:Description rdf:about="#klzagent-7293-virtMem">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Virtual Memory Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">19.85</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#VirtualMemoryUsed"/>
</rdf:Description>
<rdf:Description rdf:about="#klzagent-20399-virtMem">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Virtual Memory Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">19.78</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#VirtualMemoryUsed"/>
</rdf:Description>
<rdf:Description rdf:about="#klzagent-24848-realMem">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Real Memory Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">3.23</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#RealMemoryUsed"/>
</rdf:Description>
<rdf:Description rdf:about="#klzagent-27517-realMem">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Real Memory Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">3.17</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#RealMemoryUsed"/>
</rdf:Description>
<rdf:Description rdf:about="#klzagent-7293-realMem">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Real Memory Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">3.03</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#RealMemoryUsed"/>
</rdf:Description>
<rdf:Description rdf:about="#klzagent-3430-realMem">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Real Memory Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">3.02</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#RealMemoryUsed"/>
</rdf:Description>
<rdf:Description rdf:about="#klzagent-20399-realMem">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Real Memory Utilization</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">3.01</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#RealMemoryUsed"/>
</rdf:Description>
<rdf:Description rdf:about="#diskUsed-/">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Percentage space used</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">15</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#DiskSpaceUsed"/>
</rdf:Description>
<rdf:Description rdf:about="#diskUsed-/dev">
    <rdf:type rdf:resource="http://open-services.net/ns/ems#Measure"/>
    <dcterms:title>Percentage space used</dcterms:title>
    <ems:numericValue rdf:datatype="http://www.w3.org/2001/XMLSchema#double">1</ems:numericValue>
    <ems:unitOfMeasure rdf:resource="http://dbpedia.org/resource/Percentage"/>
    <ems:metric rdf:resource="http://open-services.net/ns/perfmon#DiskSpaceUsed"/>
</rdf:Description>
</rdf:RDF>

```

Documentation library

This appendix contains information about the publications related to IBM Tivoli Monitoring and to the commonly shared components of Tivoli Management Services.

These publications are listed in the following categories:

- IBM Tivoli Monitoring library
- Related publications

For information about accessing and using the publications, select **Using the publications** in the **Contents** pane of the IBM Tivoli Monitoring and OMEGAMON XE Information Center at <http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/index.jsp>.

To find a list of new and changed publications, click **What's new** on the Welcome page of the IBM Tivoli Monitoring and OMEGAMON XE Information Center. To find publications from the previous version of a product, click **Previous versions** under the name of the product in the **Contents** pane.

IBM Tivoli Monitoring library

The following publications provide information about IBM Tivoli Monitoring and about the commonly shared components of Tivoli Management Services:

- *Quick Start Guide*
Introduces the components of IBM Tivoli Monitoring.
- *Installation and Setup Guide, SC22-5445*
Provides instructions for installing and configuring IBM Tivoli Monitoring components on Windows, Linux, and UNIX systems.
- *Program Directory for IBM Tivoli Management Services on z/OS, GI11-4105*
Gives instructions for the SMP/E installation of the Tivoli Management Services components on z/OS.
- *High Availability Guide for Distributed Systems, SC22-5455*
Gives instructions for several methods of ensuring the availability of the IBM Tivoli Monitoring components.
- *IBM Tivoli zEnterprise Monitoring Agent Installation and Configuration Guide, SC14-7358*
Provides instructions for installing and configuring Tivoli zEnterprise monitoring agent components on Windows, Linux, and UNIX systems. Also includes migration and backup information, Enterprise Common Collector troubleshooting, Hardware Management Console configuration, and use of the command line interface or APIs to customize the collector. This guide complements the *Tivoli zEnterprise Monitoring Agent User's Guide*.
- *Administrator's Guide, SC22-5446*
Describes the support tasks and functions required for the Tivoli Enterprise Portal Server and clients, including Tivoli Enterprise Portal user administration.
- *Command Reference, SC22-5448*
Provides detailed syntax and parameter information, as well as examples, for the commands you can use in IBM Tivoli Monitoring.

- *Messages, SC22-5450*
Lists and explains messages generated by all IBM Tivoli Monitoring components and by z/OS-based Tivoli Management Services components (such as Tivoli Enterprise Monitoring Server on z/OS and TMS:Engine).
- *Troubleshooting Guide, GC22-5449*
Provides information to help you troubleshoot problems with the software.
- Tivoli Enterprise Portal online help
Provides context-sensitive reference information about all features and customization options of the Tivoli Enterprise Portal. Also gives instructions for using and administering the Tivoli Enterprise Portal.
- *Tivoli Enterprise Portal User's Guide, SC22-5447*
Complements the Tivoli Enterprise Portal online help. The guide provides hands-on lessons and detailed instructions for all Tivoli Enterprise Portal features.
- *Agent Builder User's Guide, SC32-1921*
Explains how to use the Agent Builder for creating monitoring agents and their installation packages, and for adding functions to existing agents.
- *Performance Analyzer User's Guide, SC27-4004*
Explains how to use the Performance Analyzer to understand resource consumption trends, identify problems, resolve problems more quickly, and predict and avoid future problems.
- *IBM Tivoli zEnterprise Monitoring Agent User's Guide, SC14-7359*
Complements the Tivoli zEnterprise monitoring agent online help. The guide provides reference information about the interface, usage scenarios, agent troubleshooting information, and information about Tivoli Common Reporting reports. This guide complements the *Tivoli zEnterprise Monitoring Agent Installation and Configuration Guide*.

Documentation for the base agents

If you purchased IBM Tivoli Monitoring as a product, you received a set of base monitoring agents as part of the product. If you purchased a monitoring agent product (for example, an OMEGAMON XE product) that includes the commonly shared components of Tivoli Management Services, you did not receive the base agents.

The following publications provide information about using the base agents.

- Operating system agents:
 - *Windows OS Agent User's Guide, SC22-5451*
 - *UNIX OS Agent User's Guide, SC22-5452*
 - *Linux OS Agent User's Guide, SC22-5453*
 - *IBM i Agent User's Guide, SC22-5454*
- Agentless operating system monitors:
 - *Agentless Monitoring for Windows Operating Systems User's Guide, SC23-9765*
 - *Agentless Monitoring for AIX Operating Systems User's Guide, SC23-9761*
 - *Agentless Monitoring for HP-UX Operating Systems User's Guide, SC23-9763*
 - *Agentless Monitoring for Solaris Operating Systems User's Guide, SC23-9764*
 - *Agentless Monitoring for Linux Operating Systems User's Guide, SC23-9762*
- Warehouse agents:
 - *Warehouse Summarization and Pruning Agent User's Guide, SC22-5457*

- *Warehouse Proxy Agent User's Guide*, SC22-5456
- System P agents:
 - *AIX Premium Agent User's Guide*, SA23-2237
 - *CEC Base Agent User's Guide*, SC23-5239
 - *HMC Base Agent User's Guide*, SA23-2239
 - *VIOS Premium Agent User's Guide*, SA23-2238
- Other base agents:
 - *Tivoli Log File Agent User's Guide*, SC14-7484
 - *Systems Director base Agent User's Guide*, SC27-2872

Related publications

For information about related products and publications select **OMEGAMON XE shared publications** or other entries in the **Contents** pane of the IBM Tivoli Monitoring and OMEGAMON XE Information Center at <http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/index.jsp> .

Other sources of documentation

You can also obtain technical documentation about IBM Tivoli Monitoring and related products from the following sources:

- Service Management Connect (SMC)

For introductory information about SMC, see IBM Service Management Connect (<http://www.ibm.com/developerworks/servicemanagement>).

For information about Tivoli products, see the Application Performance Management community on SMC at IBM Service Management Connect > Application Performance Management (<http://www.ibm.com/developerworks/servicemanagement/apm>).

Connect, learn, and share with Service Management professionals. Get access to developers and product support technical experts who provide their perspectives and expertise. Using SMC, you can:

- Become involved with transparent development, an ongoing, open engagement between external users and developers of Tivoli products where you can access early designs, sprint demos, product roadmaps, and pre-release code.
- Connect one-on-one with the experts to collaborate and network about Tivoli and Integrated Service Management.
- Benefit from the expertise and experience of others using blogs.
- Collaborate with the broader user community using wikis and forums.

- Tivoli wikis

IBM Service Management Connect > Application Performance Management (<http://www.ibm.com/developerworks/servicemanagement/apm>) includes a list of relevant Tivoli wikis that offer best practices and scenarios for using Tivoli products, white papers contributed by IBM employees, and content created by customers and business partners.

Two of these wikis are of particular relevance to IBM Tivoli Monitoring:

- The IBM Tivoli Monitoring Wiki (<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/Tivoli%20Monitoring/page/Home>) provides information about IBM Tivoli Monitoring and related distributed products, including IBM Tivoli Composite Application Management products.

- The Tivoli System z[®] Monitoring and Application Management Wiki provides information about the OMEGAMON XE products, NetView[®] for z/OS, Tivoli Monitoring Agent for z/TPF, and other System z monitoring and application management products.
- IBM Integrated Service Management Library
<http://www.ibm.com/software/brandcatalog/ismlibrary/>
IBM Integrated Service Management Library is an online catalog that contains integration documentation and other downloadable product extensions.
- Redbooks[®]
<http://www.redbooks.ibm.com/>
IBM Redbooks and Redpapers include information about products from platform and solution perspectives.
- Technotes
Technotes provide the latest information about known product limitations and workarounds. You can find Technotes through the IBM Software Support Web site at <http://www.ibm.com/software/support/>.

Support information

If you have a problem with your IBM software, you want to resolve it quickly. IBM provides ways for you to obtain the support you need.

Online

The following sites contain troubleshooting information:

- Go to the IBM Support Portal (<http://www.ibm.com/support/entry/portal/software>) and follow the instructions.
- Go to IBM Service Management Connect > Application Performance Management (<http://www.ibm.com/developerworks/servicemanagement/apm>) and select the appropriate wiki.

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 IBM Support Assistant (<http://www-01.ibm.com/software/support/isa>).

Troubleshooting Guide

For more information about resolving problems, see the product's Troubleshooting Guide.

Using IBM Support Assistant

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

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

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

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

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

1. Start the IBM Support Assistant application.
2. Select **Updater** on the Welcome page.
3. Select **New Properties and Tools** or select the **New Plug-ins** tab (depending on the version of IBM Support Assistant installed).
4. Under **Tivoli**, select your product, and then click **Install**. Be sure to read the license and description.

If your product is not included on the list under **Tivoli**, no plug-in is available yet for the product.

5. Read the license and description, and click **I agree**.
6. Restart the IBM Support Assistant.

Obtaining fixes

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

1. Go to the IBM Software Support website at <http://www.ibm.com/software/support>.
2. Under **Select a brand and/or product**, select **Tivoli**.
If you click **Go**, the **Search within all of Tivoli support** section is displayed. If you don't click **Go**, you see the **Select a product** section.
3. Select your product and click **Go**.
4. Under **Download**, click the name of a fix to read its description and, optionally, to download it.

If there is no **Download** heading for your product, supply a search term, error code, or APAR number in the field provided under **Search Support (this product)**, and click **Search**.

For more information about the types of fixes that are available, see the *IBM Software Support Handbook* at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

Receiving weekly support updates

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

1. Go to the IBM Software Support website at <http://www.ibm.com/software/support>.
2. Click **My support** in the far upper-right corner of the page under **Personalized support**.
3. If you have already registered for **My support**, sign in and skip to the next step. If you have not registered, click **register now**. Complete the registration form using your e-mail address as your IBM ID and click **Submit**.
4. The **Edit profile** tab is displayed.
5. In the first list under **Products**, select **Software**. In the second list, select a product category (for example, **Systems and Asset Management**). In the third list, select a product sub-category (for example, **Application Performance & Availability** or **Systems Performance**). A list of applicable products is displayed.
6. Select the products for which you want to receive updates.
7. Click **Add products**.
8. After selecting all products that are of interest to you, click **Subscribe to email** on the **Edit profile** tab.
9. In the **Documents** list, select **Software**.
10. Select **Please send these documents by weekly email**.
11. Update your e-mail address as needed.
12. Select the types of documents you want to receive.
13. Click **Update**.

If you experience problems with the **My support** feature, you can obtain help in one of the following ways:

Online

Send an e-mail message to erchelp@ca.ibm.com, describing your problem.

By phone

Call 1-800-IBM-4You (1-800-426-4968).

Contacting IBM Software Support

IBM Software Support provides assistance with product defects. The easiest way to obtain that assistance is to open a PMR or ETR directly from the IBM Support Assistant.

Before contacting IBM Software Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenance contract that you need depends on the type of product you have:

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus®, and Rational® products, as well as DB2 and WebSphere® products that run on Windows or UNIX operating systems), enroll in Passport Advantage® in one of the following ways:

Online

Go to the Passport Advantage website at http://www-306.ibm.com/software/howtobuy/passportadvantage/pao_customers.htm.

By telephone

For the telephone number to call in your country, go to the IBM Software Support website at <http://techsupport.services.ibm.com/guides/contacts.html> and click the name of your geographic region.

- For customers with Subscription and Support (S & S) contracts, go to the Software Service Request website at <https://techsupport.services.ibm.com/ssr/login>.
- For customers with Linux, iSeries, pSeries, zSeries, and other support agreements, go to the IBM Support Line website at <http://www.ibm.com/services/us/index.wss/so/its/a1000030/dt006>.
- For IBM eServer™ software products (including, but not limited to, DB2 and WebSphere products that run in zSeries, pSeries, and iSeries environments), you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for eServer software products, go to the IBM Technical Support Advantage website at <http://www.ibm.com/servers/eserver/techsupport.html>.

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. From other countries, go to the contacts page of the *IBM Software Support Handbook* on the web at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click the name of your geographic region for telephone numbers of people who provide support for your location.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law :

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to

IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© IBM 2013. Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2013. All rights reserved.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of The Minister for the Cabinet Office, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Index

A

- actions 120, 121, 122, 123
- agent
 - trace logs 209
- agents
 - instance names 7
 - troubleshooting 217
- agents, remote monitoring 1
- alerts 1
- AMS Recycle Agent Instance action 120
- AMS Reset Agent Daily Restart Count action 120
- AMS Start Agent action 121
- AMS Start Agent Instance action 121
- AMS Start Management action 122
- AMS Stop Agent action 122
- AMS Stop Management action 123
- attribute groups
 - overview 25
- attributes
 - overview 25

B

- built-in troubleshooting features 207

C

- calculate historical data disk space 103
- capacity planning for historical data 103
- code, product 3
- column manipulation 141
- commands, Take Action 119
- components 3
- configuration 5
- customer support 313

D

- dashboard
 - URL 142
- Dashboard Health Checks 127
- dashboards
 - See also* server dashboards
 - setting a trace 143
- data
 - trace logs 208
- data collection 263
- data provider logs 209
- database agent installation problems 212
- developerWorks 309
- disk capacity planning for historical data 103
- disk space requirements 6

E

- environment
 - features 1
- event
 - mapping 237

F

- features, Monitoring Agent for Linux OS 1
- files
 - agent trace 209
 - installation trace 209
 - other trace log 209
 - trace logs 208
- filters 141
- fixes, obtaining 312

G

- gathering support information 207

H

- historical data
 - calculate disk space 103
 - disk capacity planning 103

I

- IBM Redbooks 311
- IBM Software Support 207
- IBM Support Assistant 311
- IBM Tivoli Enterprise Console
 - event mapping 237
 - optional product 4
- IBM Tivoli Monitoring: Linux OS Agent
 - performance considerations 221
- information
 - troubleshooting 207
- Infrastructure Management Dashboards for Servers
 - See* server dashboards
- installation 5
 - log file 209
 - problems 212
- Integrated Service Management Library 309
- interface, user 4
 - troubleshooting for Tivoli Enterprise Portal 219
- ISA 311

L

- launch to the Tivoli Enterprise Portal 142
- limited user permissions, upgrading your warehouse
 - with 228
- Linux agent installation problems 212
- Linux_AMS_Alert_Critical situation 109
- Linux_BP_AvgCpuBusyPct1h_Critic situation 109
- Linux_BP_CpuBusyPct_Critical situation 109
- Linux_BP_CpuWaitIOPct_Warning situation 110
- Linux_BP_LoadAvg5min_Critical situation 110
- Linux_BP_NetTotalErrPct_Warning situation 110
- Linux_BP_NumberZombies_Warning situation 110
- Linux_BP_ProcHighCpu_Critical situation 110
- Linux_BP_ProcMissing_Critical situation 111
- Linux_BP_SpaceUsedPct_Critical situation 111
- Linux_BP_SpaceUsedPctCustom_Cri situation 111

- Linux_BP_SwapSpaceUsedPct_Criti situation 111
- Linux_Fragmented_File_System situation 111
- Linux_Fragmented_File_System_2 situation 112
- Linux_High_CPU_Overload situation 112
- Linux_High_CPU_Overload_2 situation 112
- Linux_High_CPU_System situation 112
- Linux_High_CPU_System_2 situation 112
- Linux_High_Packet_Collisions situation 112
- Linux_High_Packet_Collisions_2 situation 113
- Linux_High_RPC_Retransmit situation 113
- Linux_High_RPC_Retransmit_2 situation 113
- Linux_High_Zombies situation 113
- Linux_High_Zombies_2 situation 113
- Linux_Low_Pct_Inodes situation 113
- Linux_Low_Pct_Inodes_2 situation 114
- Linux_Low_percent_space situation 114
- Linux_Low_percent_space_2 situation 114
- Linux_Low_Space_Available situation 114
- Linux_Low_Space_Available_2 situation 114
- Linux_Network_Status situation 114
- Linux_Network_Status_2 situation 115
- Linux_NFS_Buffer_High situation 115
- Linux_NFS_Buffer_High_2 situation 115
- Linux_NFS_Getattr_High situation 115
- Linux_NFS_Getattr_High_2 situation 115
- Linux_NFS_rdlink_high situation 115
- Linux_NFS_rdlink_high_2 situation 116
- Linux_NFS_Read_High situation 116
- Linux_NFS_Read_High_2 situation 116
- Linux_NFS_Writes_High situation 116
- Linux_NFS_Writes_High_2 situation 116
- Linux_Packets_Error situation 116
- Linux_Packets_Error_2 situation 117
- Linux_Process_High_Cpu situation 117
- Linux_Process_High_Cpu_2 situation 117
- Linux_Process_High_Instant_CPU situation 117
- Linux_Process_stopped situation 117
- Linux_Process_stopped_2 situation 117
- Linux_RPC_Bad_Calls situation 118
- Linux_RPC_Bad_Calls_2 situation 118
- Linux_System_Thrashing situation 118
- Linux_System_Thrashing_2 situation 118
- logging
 - agent trace logs 209
 - built-in features 207
 - installation log files 209
 - location and configuration of logs 208
 - trace log files 208

M

- managed system groups
 - server dashboards 132
- memory requirements 5
- messages
 - built-in features 207
- Monitoring Agent for Linux OS
 - components 3
 - features 1
- Monitoring Agent for Linux OS installation problems 212
- monitoring agents, remote 1
- monitoring servers 1

N

- non-administrator user 7
- non-root user 7

O

- operating systems 5
- other requirements 6

P

- path names
 - for trace logs 208
- performance considerations 221
- permissions, upgrading your warehouse with limited user 228
- policies
 - list of all 125
 - overview 125
 - predefined 125
- problem resolution 311
- problems and workarounds 212
- product code 3
- purposes
 - troubleshooting 207

Q

- queries, using attributes 25

R

- Redbooks 309, 311
- remote deployment
 - troubleshooting 220
- remote monitoring agents 1
- requirements
 - disk space 6
 - memory 5
 - operating system 5
 - other 6

S

- Sample_kill_Process Take Action command 123
- server dashboards
 - background information 127
 - managed system groups overview 132
 - opening home dashboard 130
 - opening situation event results 131
 - situation event results 136
 - situation events 135
- Service Management Connect 309, 311
- situation events
 - dashboard 135
 - opening event results dashboard 131
 - results dashboard 136
- situations
 - general troubleshooting 223
 - Linux_AMS_Alert_Critical 109
 - Linux_BP_AvgCpuBusyPct1h_Critic 109
 - Linux_BP_CpuBusyPct_Critical 109
 - Linux_BP_CpuWaitIOPct_Warning 110
 - Linux_BP_LoadAvg5min_Critical 110

situations (*continued*)

- Linux_BP_NetTotalErrPct_Warning 110
- Linux_BP_NumberZombies_Warning 110
- Linux_BP_ProcHighCpu_Critical 110
- Linux_BP_ProcMissing_Critical 111
- Linux_BP_SpaceUsedPct_Critical 111
- Linux_BP_SpaceUsedPctCustom_Cri 111
- Linux_BP_SwapSpaceUsedPct_Criti 111
- Linux_Fragmented_File_System 111
- Linux_Fragmented_File_System_2 112
- Linux_High_CPU_Overload 112
- Linux_High_CPU_Overload_2 112
- Linux_High_CPU_System 112
- Linux_High_CPU_System_2 112
- Linux_High_Packet_Collisions 112
- Linux_High_Packet_Collisions_2 113
- Linux_High_RPC_Retransmit 113
- Linux_High_RPC_Retransmit_2 113
- Linux_High_Zombies 113
- Linux_High_Zombies_2 113
- Linux_Low_Pct_Inodes 113
- Linux_Low_Pct_Inodes_2 114
- Linux_Low_percent_space 114
- Linux_Low_percent_space_2 114
- Linux_Low_Space_Available 114
- Linux_Low_Space_Available_2 114
- Linux_Network_Status 114
- Linux_Network_Status_2 115
- Linux_NFS_Buffer_High 115
- Linux_NFS_Buffer_High_2 115
- Linux_NFS_Getattr_High 115
- Linux_NFS_Getattr_High_2 115
- Linux_NFS_rmlink_high 115
- Linux_NFS_rmlink_high_2 116
- Linux_NFS_Read_High 116
- Linux_NFS_Read_High_2 116
- Linux_NFS_Writes_High 116
- Linux_NFS_Writes_High_2 116
- Linux_Packets_Error 116
- Linux_Packets_Error_2 117
- Linux_Process_High_Cpu 117
- Linux_Process_High_Cpu_2 117
- Linux_Process_High_Instant_CPU 117
- Linux_Process_stopped 117
- Linux_Process_stopped_2 117
- Linux_RPC_Bad_Calls 118
- Linux_RPC_Bad_Calls_2 118
- Linux_System_Thrashing 118
- Linux_System_Thrashing_2 118
- list of all 107
- overview 107
- predefined 107
- specific troubleshooting 221

situations, using attributes 25

SMC 309, 311

Software Support 311

- contacting 313
- receiving weekly updates 312

sorting a table 141

standardization 1

support

- gathering information for 207

support assistant 311

Support Assistant 311

T

table controls 141

Take Action commands

- AMS Recycle Agent Instance 120
- AMS Reset Agent Daily Restart Count 120
- AMS Start Agent 121
- AMS Start Agent Instance 121
- AMS Start Management 122
- AMS Stop Agent 122
- AMS Stop Management 123
- overview 119
- Sample_kill_Process 123

Technotes 309

Tivoli Availability Portal

- how to use 1

Tivoli Common Reporting

- troubleshooting 220

Tivoli Data Warehouse 4

Tivoli Enterprise Console 4

Tivoli Enterprise Monitoring Server 3

Tivoli Enterprise Portal 127, 142

- component 3
- troubleshooting 219

trace

- server dashboards 143

trace logs 208

- directories 208

troubleshooting 207, 212

- agents 217
- built-in features 207
- installation 212
- installation logs 209
- remote deployment 220
- situations 221, 223
- Tivoli Common Reporting 220
- Tivoli Enterprise Portal 219
- uninstallation 212
- uninstallation logs 209

U

uninstallation

- log file 209
- problems 212

upgrading for warehouse summarization 227

upgrading your warehouse with limited user permissions 228

URL of dashboard 142

user interfaces options 4

user permissions, upgrading your warehouse with limited 228

W

Warehouse Proxy agent 4

warehouse summarization

- upgrading for
- overview 227

Warehouse Summarization and Pruning agent 4

warehouse summarization upgrading

- affected attribute groups and supporting scripts 232
- DB2 warehouse database procedure 233
- effects on summarized attributes 227
- MS SQL warehouse database procedure 234
- Oracle warehouse database procedure 233
- procedures for running scripts 232

- warehouse summarization upgrading *(continued)*
 - table summary 230
 - tables in the warehouse 227
 - types of table changes 229
 - upgrading your warehouse 231
- workarounds 212
 - agents 217
 - remote deployment 220
 - situations 221
 - Tivoli Common Reporting 220
 - Tivoli Enterprise Portal 219
- workspaces
 - list of all 11
 - overview 11
 - predefined 11



Printed in USA

SC22-5453-00

