Tivoli® Monitoring: Linux OS Agent

Version 6.2.0

User's Guide

SC32-9447-01

Tivoli® Monitoring: Linux OS Agent

Version 6.2.0

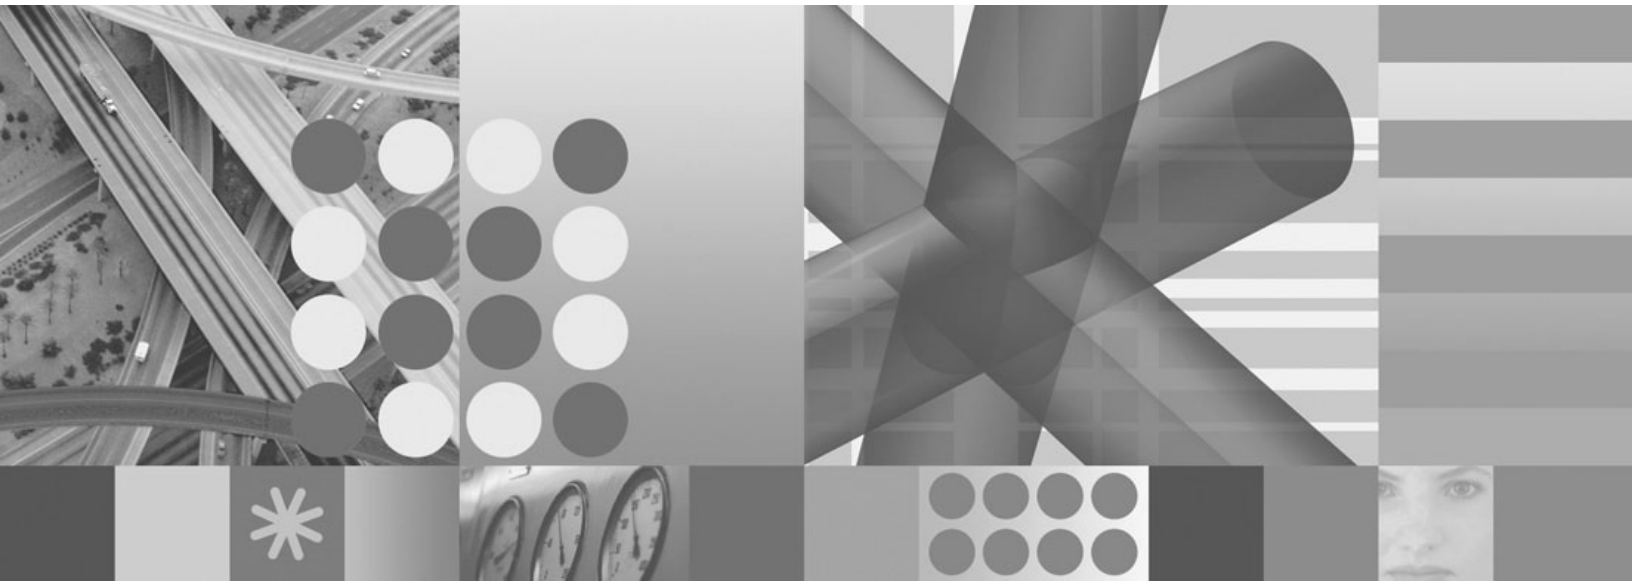

User's Guide

SC32-9447-01

Note
Before using this information and the product it supports, read the information in "Notices" on page 133.

This edition applies to version 6.2 of the IBM Tivoli Monitoring: Linux OS Agent (5724-C04) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Tables

Chapter 1. Overview of the Monitoring Agent for Linux OS

The Monitoring Agent for Linux OS provides you with the capability to monitor Linux, and to perform basic actions with Linux. This chapter provides a description of the features, components, and interface options for the Monitoring Agent for Linux OS.

IBM Tivoli Monitoring overview

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 Appendix D, “Documentation library,” on page 129 for complete information about IBM Tivoli Monitoring and the Tivoli Enterprise Portal.

Features of the Monitoring Agent for Linux OS

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 appears 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 6.2 of the Monitoring Agent for Linux OS, the following enhancements have been made:

- Additional supported operating systems as listed in Chapter 2, “Requirements for the monitoring agent,” on page 5
- Enablement of IBM® Tivoli® License Manager reporting
- New configuration option
  - Enabling history collection for the Linux® Host Availability attributes. For more information, see“Linux Host Availability Attributes” on page 43.
- New workspaces
  - File Information
- New attribute groups
  - All Users
  - File Comparison
  - File Information
  - File Pattern
  - IP Address
  - Linux Group
  - Linux Host Availability
  - Machine Information
- Updated klz.baroc file to support TEC event mapping
- Updated resource model mapping files

Note: These enhancements include ones made for the various IBM Tivoli Monitoring fix packs since the release of IBM Tivoli Monitoring 6.1.

Monitoring Agent for Linux OS components

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 that contains the 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, Monitoring Agent for Linux OS, which collects and distributes data to a Tivoli Enterprise Monitoring Server.
- 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

This chapter contains information about the following topics and procedures relevant to the installation and configuration of the Monitoring Agent for Linux OS.

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 on page 6.

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

Operating system	**Linux**
Operating system versions	Linux: • Linux on zSeries – RedHat Enterprise Linux AS 3 (31-bit or 64-bit) – RedHat Enterprise Linux AS 4 (31-bit or 64-bit) – RedHat Enterprise Linux AS 5 (31-bit or 64-bit) – SUSE Linux Enterprise Server 8 (31-bit or 64-bit) – SUSE Linux Enterprise Server 9 (31-bit or 64-bit) – SUSE Linux Enterprise Server 10 (31-bit or 64-bit) • Linux on Intel® (32-bit) – RedHat Enterprise Linux AS/ES 2.1 – RedHat Enterprise Linux AS/ES 3 – RedHat Enterprise Linux AS/ES 4 – RedHat Enterprise Linux AS/ES 5 – SUSE Linux Enterprise Server 8 – SUSE Linux Enterprise Server 9 – SUSE Linux Enterprise Server 10 – RedFlag 4.1 – Asian Linux 2.0 • Linux on pSeries – RedHat Enterprise Linux AS 4 – RedHat Enterprise Linux AS 5 – SUSE Linux Enterprise Server 9 – SUSE Linux Enterprise Server 10 • Linux on IA64 (Itanium®) – RedHat Enterprise Linux AS 4[1] – RedHat Enterprise Linux AS 5[1] – SUSE Linux Enterprise Server 9 [1] – SUSE Linux Enterprise Server 10[1] – Asian Linux 2 • Linux on x86-64 – RedHat Enterprise Linux AS 4[1] – RedHat Enterprise Linux AS 5[1] – SUSE Linux Enterprise Server 9[1] – SUSE Linux Enterprise Server 10[1] – Asian Linux 2 The Linux version must support the Korn shell (ksh) and Motif Window Manager (libmotif) for installation of the monitoring agent.
Memory	• 256 MB RAM at a minimum although 512 MB or higher for better performance
Disk space	• 100 MB of disk space for the base monitoring agent • Historical data disk space: see “Disk capacity planning for historical data” on page 70

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

Operating system	**Linux**
Other requirements	• 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 libs to be installed for the agent to work correctly. The latest versions of libstdc++, libgcc, compat-libstdc++, and libXp are required for the agent to work correctly.
Notes: 1. In native 64-bit mode, not tolerance mode.	

Note: For the most current information about the operating systems that are supported, see the following URL:

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

When you get to that site, click **Tivoli platform and database support matrix link** at the bottom of the window.

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 hostname
- Agent product node type

An agent name truncation problem can occur when the network domain name is included in the network hostname portion of the agent name. For example, instead of just the hostname myhost1 being used, the resulting hostname might be myhost1.acme.north.prod.com. Inclusion of the network domain name causes the agent name in the example above to expand to SERVER1:myhost1.acme.north.prod.com:KXX. This resulting name is 39 characters long. It is truncated to 32 characters resulting in the name SERVER1:myhost1.acme.north.prod.

The agent name truncation is only a problem if there is more than one monitoring agent on the same system. In this case, the agent name truncation can result in collisions between agent products attempting to register using the same truncated name value. When truncated agent names collide on the same system, this can lead to Tivoli Enterprise™ Monitoring Server problems with corrupted EIB tables. The agent name collision in the Tivoli Enterprise Monitoring Server might cause a registered name to be associated with the wrong product.

In general, create names that are short but meaningful within your environment. Use the following guidelines:

- Each name must be unique. One name cannot match another monitoring agent name exactly.
- Each name must begin with an alpha character.
- Do not use blanks or special characters, including $, #, and @.
- Each name must be between 2 and 32 characters in length.
- Monitoring agent naming is case-sensitive on all operating systems.

See "Unique names for monitoring components" on page 117 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.

Chapter 3. How to use a monitoring agent

After you have installed and configured a Tivoli Enterprise Monitoring Agent and the agent is running, you can begin using this agent to monitor your resources. The following sources of information are relevant to installation and configuration:

- *IBM Tivoli Monitoring Installation and Setup Guide*
- *IBM Tivoli Monitoring Command Reference*
- Chapter 2, "Requirements for the monitoring agent" in the user's guide for the agent that you are installing and configuring

This chapter provides information about how to use a monitoring agent to perform the following tasks:

- "View real-time data that the agent collects"
- "Investigate an event" on page 10
- "Recover the operation of a resource" on page 10
- "Customize your monitoring environment" on page 11
- "Monitor with custom situations that meet your requirements" on page 12
- "Collect and view historical data" on page 13

For each of these tasks, there is a list of procedures that you perform to complete the task. For the tasks, there is a cross-reference to where you can find information about performing that procedure. Information about the procedures is located in subsequent chapters of this user's guide and in the following publications:

- *IBM Tivoli Monitoring User's Guide*
- *IBM Tivoli Monitoring Administrator's Guide*

View real-time data that the agent collects

After you install, configure, and start the Tivoli Enterprise Monitoring Agent, the agent begins monitoring.

Table 2 contains a list of the procedures for viewing the real-time data that the monitoring agent collects through the predefined situations. The table also contains a cross-reference to where you can find information about each procedure.

Table 2. View real-time data

Procedure	Where to find information
View the hierarchy of your monitored resources from a system point of view (Navigator view organized by operating system type, monitoring agents, and workspaces).	*IBM Tivoli Monitoring User's Guide:* ″Navigating through workspaces″ (in ″Monitoring: real-time and event-based″ chapter)
View the indicators of real or potential problems with the monitored resources (Navigator view).	

Table 2. View real-time data (continued)

Procedure	Where to find information
View changes in the status of the resources that are being monitored (Enterprise Message Log view).	*IBM Tivoli Monitoring User's Guide:* "Using workspaces" (in "Monitoring: real-time and event-based" chapter) Chapter 4, "Workspaces reference," on page 15 in this guide
View the number of times an event has been opened for a situation during the past 24 hours (Open Situations Account view).	*IBM Tivoli Monitoring User's Guide:* "Using workspaces" (in "Monitoring: real-time and event-based" chapter) Chapter 4, "Workspaces reference," on page 15 in this guide Chapter 6, "Situations reference," on page 73 in this guide
Manipulate the views in a workspace.	*IBM Tivoli Monitoring User's Guide:* "Using views" (in "Monitoring: real-time and event-based" chapter)

Investigate an event

When the conditions of a situation have been met, an event indicator is displayed in the Navigator. When an event occurs, you want to obtain information about that event so you can correct the conditions and keep your enterprise running smoothly.

Table 3 contains a list of the procedures for investigating an event and a cross-reference to where you can find information about each procedure.

Table 3. Investigating an event

Procedure	Where to find information
Determine which situation raised the event and identify the attributes that have values that are contributing to the alert.	*IBM Tivoli Monitoring User's Guide:* "Opening the situation event workspace" (in "Monitoring: real-time and event-based" chapter, "Responding to alerts" section)
Review available advice.	Chapter 4, "Workspaces reference," on page 15 in this guide
Notify other users that you have taken ownership of the problem related to an event and are working on it.	*IBM Tivoli Monitoring User's Guide:* "Acknowledging a situation event" (in "Monitoring: real-time and event-based" chapter, "Responding to alerts" section)
Remove the event from the Navigator.	*IBM Tivoli Monitoring User's Guide:* "Closing the situation event workspace" (in "Monitoring: real-time and event-based" chapter, "Responding to alerts" section)

Recover the operation of a resource

When you find out that a resource is not operating as desired, you can control it manually or automatically using Take Action commands.

Table 4 contains a list of the procedures for recovering the operation of a resource and a cross-reference to where you can find information about each procedure.

Table 4. Recover the operation of a resource

Procedure	Where to find information
Take an action on a resource manually.	*IBM Tivoli Monitoring User's Guide:* • ″Other views″ (in ″Custom workspaces″ chapter, ″Workspace views″ section) • ″Take action: Reflex automation″ (in Situations for event-based monitoring″ chapter, ″Event-based monitoring overview″ section) • "Take action" (in "Designing customized responses" chapter) Chapter 7, “Take Action commands reference,” on page 79 in this guide
Take an action on a system condition automatically by setting up a situation to run a Take Action command.	*IBM Tivoli Monitoring User's Guide:* ″Situations for event-based monitoring″ chapter • ″Customize a situation″ • ″Create a situation″ • ″Specify an action to take″ • ″Distribute the situation″ Chapter 7, “Take Action commands reference,” on page 79 in this guide
Take multiple actions on system conditions automatically using a policy.	*IBM Tivoli Monitoring User's Guide:* ″Policies for automation″ chapter • ″Creating a policy″ • ″Maintaining policies″ • ″Workflows window″ Chapter 8, “Policies reference,” on page 81 in this guide
Take actions across systems, agents, or computers using a policy.	

Customize your monitoring environment

You can change how your monitoring environment looks by creating new workspaces with one or more views in it.

Table 5 contains a list of the procedures for customizing your monitoring environment and a cross-reference to where you can find information about each procedure.

Table 5. Customizing your monitoring environment

Procedure	Where to find information
Display data in tables or charts (views) in the Tivoli Enterprise Portal.	*IBM Tivoli Monitoring User's Guide:* • ″Custom workspaces″ • ″Table and chart views″

Table 5. Customizing your monitoring environment (continued)

Procedure	Where to find information
Display an overview of changes in the status of situations for your monitored resources (Message Log View).	*IBM Tivoli Monitoring User's Guide:* "Message log view" (in "Situation event views: message log, situation event console and graphic" chapter)
Specify which attributes to retrieve for a table or chart so you can retrieve only the data you want by creating custom queries.	*IBM Tivoli Monitoring User's Guide:* "Creating custom queries" (in "Table and chart views" chapter) Chapter 5, "Attributes reference," on page 27 in this guide
Build links from one workspace to another.	*IBM Tivoli Monitoring User's Guide:* • "Link from a workspace" (in "Custom workspaces" chapter) • "Link from a table or chart" (in "Table and chart views" chapter)
Identify which predefined situations started running automatically when you started the Tivoli Enterprise Monitoring Server.	*IBM Tivoli Monitoring User's Guide:* "What the enterprise workspace shows" (in "Monitoring: real-time and event-based" chapter, "Using workspaces" section) Chapter 6, "Situations reference," on page 73 in this guide
Determine whether to run situations as defined, modify the values in situations, or create new situations to detect possible problems.	Chapter 6, "Situations reference," on page 73 in this guide

Monitor with custom situations that meet your requirements

When your environment requires situations with values that are different from those in the existing situations, or when you need to monitor conditions not defined by the existing situations, you can create custom situations to detect problems with resources by creating an entirely new situation.

You can specify the following information for a situation:

- Name
- Attribute group and attributes
- Qualification to evaluate multiple rows when a situation has a multiple-row attribute group (display item)
- Formula
- Take Action commands
- Run at startup
- Sampling interval
- Persistence
- Manual or automatic start
- Severity
- Clearing conditions
- Expert Advice
- When a true situation closes

- Available Managed Systems
- Whether to send a Tivoli Enterprise Console event
- Event severity

Table 6 contains a list of the procedures for monitoring your resources with custom situations that meet your requirements and a cross-reference to where you can find information about each procedure.

Table 6. Monitor with custom situations

Procedure	Where to find information
Create an entirely new situation.	*IBM Tivoli Monitoring User's Guide:* "Creating a new situation" (in "Situations for event-based monitoring" chapter, "Creating a situation" section) Chapter 5, "Attributes reference," on page 27 in this guide
Run a situation on a managed system.	*IBM Tivoli Monitoring User's Guide:* "Situations for event-based monitoring" chapter • "Associating situations with navigator items" • "Distribute the situation" (in "Customizing a situation" section) • "Starting, stopping or deleting a situation"

Collect and view historical data

When you collect historical data, you specify the following configuration requirements:

- Attribute groups for which to collect data
- Collection interval
- Summarization and pruning of attribute groups
- Roll-off interval to a data warehouse, if any
- Where to store the collected data (at the agent or the Tivoli Enterprise Management Server)

Table 7 on page 14 contains a list of the procedures for collecting and viewing historical data and a cross-reference to where you can find information about each procedure.

Table 7. Collect and view historical data

Procedure	Where to find information
Configure and start collecting short-term data (24 hours).	*IBM Tivoli Monitoring User's Guide:* "Historical reporting" (in "Table and chart views" chapter) *IBM Tivoli Monitoring Administrator's Guide* "Disk capacity planning for historical data" on page 70 in this guide
Configure and start collecting longer-term data (more than 24 hours).	
View historical data in the Tivoli Enterprise Portal.	
Create reports from historical data using third-party reporting tools.	
Filter out unwanted data to see specific areas of interest.	

Chapter 4. Workspaces reference

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

About workspaces

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

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

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

More information about workspaces

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

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

Predefined workspaces

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

  - "Disk I/O Extended Rate workspace" on page 17
  - "Virtual Memory Statistics workspace" on page 24
- "Users Workspace" on page 24

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 of this chapter contain descriptions of each of these predefined workspaces, which are organized alphabetically within the group.

All Files workspace

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

- Network Activity (maximum 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 five 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 five 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 five 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

Network workspace

The Network workspace reflects the health of the network components within your monitored systems. 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 workspace

The NFS workspace is reached by right-clicking the Network workspace in the Tivoli Enterprise Portal. The NFS 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. 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 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 percent of CPU time spent in kernel mode and spent in user mode by process. The Process + Child CPU Percent Usage chart displays the cumulative percent 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 User Information workspace

The Process User Information workspace is reached by right-clicking the Process workspace in the Tivoli Enterprise Portal. The Process User Information workspace identifies process owners of your monitored Linux system and details their usage. 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 percent of CPU time spent in kernel mode and spent in user mode by process. The Process + Child CPU Percent Usage chart displays the cumulative percent 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 workspace

The RPC workspace is reached by right-clicking the Network workspace 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. 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 workspace in the Tivoli Enterprise Portal. The Sockets Information workspace reflects the health of the socket connections within your monitored systems. 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 is reached by right-clicking the File Information workspace in the Tivoli Enterprise Portal. The Specific File Information workspace contains file information. 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. It contains three views:

- CPU Usage (bar chart)
- Processor Configuration Information (table view)
- OS Version 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 is comprised of four views. The views are:

- CPU Usage (bar chart)
- Paging (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, and user nice 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 fifteen minutes. The paging 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.

Users Workspace

The Users workspace identifies logged in users. 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 into the monitored system during the monitoring period. The User Login Information table lists users, their login 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 workspace in the Tivoli Enterprise Portal. The Virtual Memory Statistics workspace provides a snapshot of your monitored systems memory usage. 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 percent 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 workspace in the Tivoli Enterprise Portal. 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 5. Attributes reference

This chapter contains information about the following topics:

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

About attributes

Attributes are the application properties being measured and reported by the Monitoring Agent for Linux OS, such as the amount of memory usage or the message ID. Some agents have fewer than 100 attributes, while others have over 1000.

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

- Chart or table views

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

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

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

More information about attributes

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

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

Attribute groups and attributes for the Monitoring Agent for Linux OS

The following sections contain descriptions of these attribute groups, which are listed alphabetically. Each description contains a list of attributes in the attribute group.

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.

Note: 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.

All Users Group Attributes

The All Users Group 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 valid values are False and True.

Name The full name of a user.

No Password True if no password is assigned to the user. The valid values are Unknown, False, and True.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

User ID The numeric ID the system assigned to a user.

User Sessions The number of login sessions this user currently has established.

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. Valid entry is an integer. Valid entry is an integer in the range 0 to 100.

CPU ID The processor ID. Valid entry is an integer 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.

Idle CPU (Percent) Percent of idle CPU time during the sampling period. Valid entry is an integer 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 percent is larger (greater than 30%), the system might be compute-bound or in a loop.

IO Wait (Percent) The percentage of time the CPU was in a wait input/output state. Valid entry is an integer in the range of 0 to 100.

System CPU (Percent) Percent of system CPU time during the sampling period. Valid entry is an integer in the range 0 to 100. Use this attribute to determine the percent 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.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

User CPU (Percent) Percent of user CPU time during the sampling period. Valid entry is an integer in the range 0 to 100. Use this attribute to determine the percent 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.

User Nice CPU (Percent) Percent of user nice CPU time during the sampling period. Valid entry is an integer in the range 0 to 100.

User to System CPU (Percent) Of the total CPU time, the percentage consumed by users. The range of possible values for this attribute is -10000 to + 10000.

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. Valid entry is an integer. Note: -1 indicates N/A.

Idle CPU Moving Average (Percent) The moving average of the idle CPU time for the system, expressed as a percent. Valid entries: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly.

Idle CPU (Percent) The current average of the idle CPU time for the system, expressed as a percent. Valid entries: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly.

System CPU Current Average (Percent) The current average of the System CPU time, expressed as a percent. Valid entries: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly.

System CPU Moving Average (Percent) The moving average of the System CPU time, expressed as a percent. Valid entries: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Total CPU Used Current Average (Percent) The current average of CPU usage, expressed as a percent. Valid entries: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly.

Total CPU Used Moving Average (Percent) The moving average of CPU usage, expressed as a percent. Valid entries: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly.

User CPU Current Average (Percent) The current average of the user CPU time, expressed as a percent. Valid entries: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly.

User CPU Moving Average (Percent) The moving average of the user CPU time, expressed as a percent. Valid entries: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly.

User Nice CPU Current Average (Percent) The current average of the user nice CPU time, expressed as a percent. Valid entries: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly.

User Nice CPU Moving Average (Percent) The moving average of the user nice CPU time, expressed as a percent. Valid entries: integers between 0 and 100, such as 85 for 85%. This average is calculated hourly.

Wait CPU Moving Average (Percent) The moving current average of the wait CPU time, expressed as a percentage. Valid entry is an integer between 0 and 100. This average is calculated hourly.

Wait CPU (Percent) The current average of the wait CPU time, expressed as a percentage. Valid entry is an integer between 0 and 100. This average is calculated hourly.

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: -1 indicates Not_Available and -2 indicates Not_Collected.

Processor Clock Speed (MHz) The processor clock speed (MHz). Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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 process model number. 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:KLZ or deux.raleigh.ibm.com:KLZ.

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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.

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

Disk Name The name of the physical disk partition where the filesystem is mounted. This is the physical location of the disk. Valid entry is an alphanumeric text string, with a maximum length of 32 characters.

File System Type The file system type, such as hsfs, nfs, tmpfs, and ufs. Valid entries are up to eight letters or numbers.

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

Inodes Free The number of inodes currently available on your filesystem. 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 and -2 indicates Not Collected.

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

Inodes Used Percent The percent of inodes currently allocated to files, calculated by dividing the Inodes Used value by the Total Inodes value. Valid entries: integers between 0 and 100, such as 85 for 85%. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

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

Size (MB) The total size of a filesystem, expressed in megabytes. For example, 1000 represents one gigabyte. Valid entry is an integer of up to 99999999. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Space Available (MB) The amount of unused space currently available to non-superusers on a filesystem, expressed in megabytes. For example, 40000 represents 40 megabytes. Valid entry is an integer of up to 99999999.

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 filesystems, relative to the disk size, you might need to evaluate reconfiguring the filesystem to distribute the files more evenly across disks.

Space Available Percent The percent of space available. Valid entry is an integer between 0 and 100, such as 10 for 10%. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Space Used (MB) The amount of disk space currently in use on a filesystem, expressed in megabytes. For example, 5000 represents five gigabytes. Valid entry is an integer of up to 99999999. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Space Used Percent The space currently used on the file system, expressed as a percent of the sum of used and available space. The Space Used Percent reflects the percent of disk space which is available to non-superusers. A high value in this column alerts you to critical disk space conditions. Valid entries: integers between 0 and 100, such as 80 for 80%. 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:KLZ or deux.raleigh.ibm.com:KLZ.

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Total Inodes The number of inodes allocated on a filesystem. For example, a value of 163817 indicates that the number of inodes allocated is 163,817. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Use this attribute when a filesystem 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 I/O Attributes

The Disk I/O 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.

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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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. Valid entry is an integer.

Blocks Read The total number of blocks read. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Blocks Written The total number of blocks written. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Device Major Number Major number of the device. Valid entry is an integer.

Device Minor Number Distinctive minor number for device. Valid entry is an integer.

Device Name Name of the device as appears under the /dev directory.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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 filesystem is mounted. This is the physical location of the disk. Valid entry is an alphanumeric text string, 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. Valid entry is an integer. Note: -1 indicates N/A.

Days Until Full Disk Moving Avg The number of days until the disk is full based on the moving average rate of disk usage. Valid entry is an integer. Note: -1 indicates N/A.

Days Until Full Disk Peak Rate Days until full disk based on the Peak Rate. Note: -1 indicates N/A.

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. Valid entry is an integer. Note: -1 indicates N/A.

Disk Usage Moving Avg (Bytes/Hr) The bytes/hour of disk usage averaged over all previous samples. Valid entry is an integer.

Disk Usage Rate (Bytes/Hr) The bytes/hour of disk usage over the last hour. Valid entry is an integer.

High Water Mark Disk Usage Rate (Bytes/Hr) The bytes/hour rate that represents the highwater mark of disk usage. Valid entry is an integer.

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

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Space Available (MB) The amount of unused space currently available to non-superusers on a filesystem, expressed in megabytes. For example, 40,000 represents 40 megabytes. Valid entry is an integer.

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 filesystems, relative to the disk size, you might need to evaluate reconfiguring the filesystem to distribute the files more evenly across disks.

Space Used (MB) The amount of disk space currently in use on a filesystem, expressed in megabytes. Valid entries For example, 5000 represents five gigabytes. Valid entry is an integer.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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 valid values include: Plain, Ignore_Whitespace, Ignore_Case, Ignore_Case_Whitespace, and Binary. The default is Plain.

File Compare Result The result of the file comparison between File_Name_1 and File_Name_2. Valid values include Same and Different. 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:KLZ or deux.raleigh.ibm.com:KLZ.

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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, MD5, and SHA1. 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. Valid values include No, Yes, and Not Available.

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.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Type The type of file. Possible values are Dir (= directory), File (= file), Sock (= socket), Link (= Link), Spec (= Special file).

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 valid values include: Normal, Ignore_Case, Inverse_Search, and Match_Whole_Words.

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:KLZ or deux.raleigh.ibm.com:KLZ.

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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 valid values include False and True.

Group ID The ID of this group.

Group Name The name of the group.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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. Valid entry is an integer.

Average Request Size (Sectors) The average size (in sectors) of the requests that were issued to the device. Valid entry is an integer.

Average Service time (ms) The average service time (in milliseconds) for I/O requests that were issued to the device. Valid entry is an integer.

Average Wait Time (ms) The average time (in milliseconds) for I/O requests issued to the device to be served. Valid entry is an integer.

Bytes Transferred Per Second The number of bytes transferred per second.

Device Name Name of the device as appears under the /dev directory. Valid entry is an alphanumeric text string, 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.

Read Requests Per Second The number of read requests that were issued, per second, to the device. Valid entry is an integer.

Read Requests Merged Per Second The number of read requests merged, per second, that were issued to the device. Valid entry is an integer.

Read Sectors Per Second The number of sectors read, per second, from the device. Valid entry is an integer.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Write Bytes Per Second The number of bytes written to the device per second.

Write Requests Per Second The number of write requests that were issued, per second, to the device. Valid entry is an integer.

Write Requests Merged Per Second The number of write requests merged that were issued, per second, to the device. Valid entry is an integer.

Write Sectors Per Second The number of sectors written to the device, per second. Valid entry is an integer.

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. Valid entry is an alphanumeric text string, 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. Valid entry is an alphanumeric text string, with a maximum length of 46 characters.

Network Interface Name The name of the network interface. Valid entry is an alphanumeric text string, with a maximum length of 32 characters.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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. Valid values include: Successful, Unsuccessful, and Error.

Server Response Time Ping operation response time in milliseconds.

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

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

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 hostname or IP Address of the target of the ping operation. Valid entry is an alphanumeric text string, with a maximum length of 128 characters.

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

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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 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. Hardware Brand will report as "IBM." Hardware Model will report as "zSeries," and the remaining metrics 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 machine. Note: the value unknown = UNKNOWN.

Number of Processors Configured The number of processors configured for this machine. 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 machine. 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 hostname for the machine. Note: the value unknown = UNKNOWN.

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

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

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.

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. Valid entry is an integer in the range 0 to 2147483647.

Bytes Transmitted Per Second The number of bytes transmitted per second by the interface. Valid entry is an integer in the range 0 to 2147483647.

Carrier Losses The number of carrier losses that occurred in the interface. Valid entry is an integer.

Collisions Percent Of the total number of packets transmitted in this sample period, the percentage involved in a collision. Valid entry is an integer.

Collisions Per Minute The number of times a packet collided with another packet per minute. Valid entry is an integer.

Device Type The device type. Valid Values include: NETROM, ETHER, EETHER, AS25, PRONET, CHAOS, IEEE802, 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. Valid entry is an integer.

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 0 to 100.

Input Errors Per Minute The number of packets with errors received per minute by the interface. Valid entry is an integer.

Input FIFO Buffer Overruns The number of input FIFO buffer overruns that occurred during the sampling period. Valid entry is an integer.

Input Packets Dropped The number of input packets dropped by the device driver. Valid entry is an integer.

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. Valid entries: Internet protocol addresses in the form a.b.c.d. where a, b, c, and d are integers in the range 0 to 255.

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

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

UP	Indicates the interface is in service
DOWN	Indicates the interface is not in service
UP_NOT_RUNNING	Indicates the interface is in service but not running
UNKNOWN	Indicates the interface is in unknown

These values are case-sensitive.

Example:**UP** means an interface is in service.

Maximum Transmission Unit The maximum packet size (in bytes) for the specified network interface. This is a fixed value. Valid entry is an integer 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. Valid entries are simple text string, alphanumeric comprised of "Interface Name, Unit Number" where:

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

Output Errors The number of packet transmission errors by the network interface. Valid entries are integers in the range 0 to 100.

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. Valid entry is an integer.

Output FIFO Buffer Overruns The number of output FIFO buffer overruns that occurred during the sampling period. Valid entry is an integer.

Output Packets Dropped The number of output packets dropped by the device driver. Valid entry is an integer.

Packet Framing Errors The number of packet framing errors that occurred in the interface. Valid entry is an integer.

Packets Received The number of packets received by the interface during the sampling period. Valid entry is an integer in the range 0 to 99999999.

Packets Received Per Second The number of packets received per second by the interface. Valid entry is an integer in the range 0 to 2147483647.

Packets Transmitted The number of packets transmitted by the interface during the sampling period. Valid entry is an integer in the range 0 to 99999999.

Packets Transmitted Per Second The number of packets transmitted per second by the interface. Valid entry is an integer in the range 0 to 2147483647.

Received Count (KB) The number of kilobytes received since the network interface was configured. Valid entry is an integer in the range 0 to 2147483647.

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:KLZ or deux.raleigh.ibm.com:KLZ.

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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. Valid entry is an integer in the range 0 to 100. 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 0 to 2147483647.

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.

Access Calls The number of access calls made to the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Access Calls Percent Of the total number of calls made to the NFS server, the percentage that were access calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Commit Calls The number of file commit calls made to the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Commit Calls Percent Of the total number of calls made to the NFS server, the percentage that were file commit calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

File Creates The number of file create calls made to the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

File Creates Percent Of the total number of calls made to the NFS server, the percentage that contained file creation operations. Valid entry is an integer 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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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. Valid entry is an integer. 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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Getattr Calls The number of calls made to the NFS server which contained a get attribute (getattr) operation. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Getattr Calls Pct Of the total number of calls made to the NFS server, the percentage that contained get attribute (getattr) operations. Valid entry is an integer 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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Link Calls Percent Of the total number of calls made to the NFS server, the percentage that were link calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Location The location of the origin of the call in the Network File System. Valid entry is an integer. A value of 0 indicates unknown, the value of 1 represents the server, and a value of 2 represents the client.

Lookups The number of lookups made on the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Lookups Percent Of the total number of calls made to the NFS server, the percentage that were lookups. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Make Directory Calls The number of make directory calls made to the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Make Directory Calls Percent Of the total number of calls made to the NFS server, the percentage that were make directory calls. Valid entry is an integer. 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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Make Node Calls Percent Of the total number of calls made to the NFS server, the percentage that were make node (mknod) calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

NFS Calls The total NFS server or client calls. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

NFS Version The software version associated with the NFS server. Valid entry is an integer. 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.

Null Calls The number of calls made to the NFS server from NFS clients which contained no data. Valid entry is an integer in the range of 0 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Null Call Percentage Of the total number of calls made to the NFS server, the percentage that contained no data. Valid entry is an integer. 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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Read Calls The number of read calls made to the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Read Calls Percent Of the total number of calls made to the NFS server, the percentage that were read calls. Valid entry is an integer. 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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Read Directory Calls Percent Of the total number of calls made to the NFS server, the percentage that were read directory calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Read Link Calls The number of read link calls made to the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Read Link Calls Percent Of the total number of calls made to the NFS server, the percentage that were read link calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Remove Directory Calls The number of remove directory calls made to the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Remove Directory Calls Percent Of the total number of calls made to the NFS server, the percentage that were remove directory calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Remove File Calls The number of file removal calls made to the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Remove File Calls Percent Of the total number of calls made to the NFS server, the percentage that were file removal calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Rename File Calls The number of file rename calls made to the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Rename File Calls Percent Of the total number of calls made to the NFS server, the percentage that were file rename calls. Valid entry is an integer. 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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Root Calls Percent Of the total number of calls made to the NFS server, the percentage that were root calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Setattr Calls The number of calls made to the NFS server which contained a set attribute (setattr) operation. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Setattr Calls Percent Of the total number of calls made to the NFS server, the percentage that contained a set attribute (setattr) operation. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Symbolic Link Calls The total number of symbolic link calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Symbolic Link Calls Percentage Of the total number of calls made to the NFS server, the percentage that were symbol link calls. Valid entry is an integer. 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:KLZ or deux.raleigh.ibm.com:KLZ.

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Write Cache Calls The number of write cache calls made to the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Write Cache Calls Percent Of the total number of calls made to the NFS server, the percentage that were write cache calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Writes The number of write calls made to the NFS server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Writes Percent Of the total number of calls made to the NFS server, the percentage that were write calls. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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:KLZ or deux.raleigh.ibm.com:KLZ.

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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.

Command Line The process command line string. Valid entry is an alphanumeric text string, with a maximum length of 256 characters.

Command Line (Unicode) The process command line string. Valid entry is a text string, with a maximum length of 512 bytes. This attribute is globalized (Unicode).

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

Cumulative Process System CPU (Percent) The percent of cumulative CPU time spent in kernel mode by process. Valid entry is an integer between 0 and 100.

Cumulative Process User CPU (Percent) The percent of cumulative CPU time spent in user mode by process. Valid entry is an integer between 0 and 100.

Data Set Resident Set (Pages) The size of the data set based on the number of pages. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

Data Size (KB) The data size (in kilobytes) of the virtual memory. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648. 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. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

Executable Size (KB) The executable size (in kilobytes) of the virtual memory. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

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. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

Kernal Priority The kernel scheduling priority (0 represents the highest priority). Valid entry is an integer between 100 and 0.

Library Size (KB) The library size (in kilobytes) of the virtual memory. This measurement represents all pages, including unused. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648. 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. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

Nice Value The standard Linux nice level (-20 represents the highest level). Valid entry is an integer 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: -1 indicates Not_Available and -2 indicates Not_Collected.

Parent Process ID The identifier for the parent process. Valid entry is an integer between 0 and 999.

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

Process Command Name The name of the process command. Valid entry is an alphanumeric text string, with a maximum length of 32 characters.

Process Command Name (Unicode) The name of the process command. Valid entry is a text string, with a maximum length of 64 bytes. This attribute is globalized (Unicode).

Process CPU ID The ID of the process CPU. Valid entry is an integer. Note: -1 indicates N/A.

Process Dirty Pages Pages that have been modified (dirty) in buffer (main memory), but not yet copied to the cache. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

Process ID The identifier of the process. Valid entry is an integer between 0 and 999.

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 kernal 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.

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.

Process State The state of the process (Sleeping, Disk, Running, Zombie, Trace, Dead, or N/A). Valid entry is an integer between -1 and 5, where:

0 = Sleeping

1 = Disk

2 = Running

3 = Zombie

4 = Trace

5 = Dead

-1 = Not_Available

Process System CPU (Percent) The percent of CPU time spent in kernel mode by process. Valid entry is an integer between 0 and 100.

Process User CPU (Percent) The percent of CPU time spent in user mode by process. Valid entry is an integer between 0 and 100.

Resident Set Size (Pages) The number of pages the process has in real memory. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

Session ID The session ID. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

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. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

Stack Size (KB) The stack size (in kilobytes) of the virtual memory. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Stack Size (MB) The stack size (in megabytes) of the virtual memory. Valid entry is an integer. This attribute displays as a floating point with a scale of 1. For example 5.2. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

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

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

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. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

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

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Total Major Faults The total number of major page faults (including child processes) since the start of the process. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

Total Minor Faults The total number of minor page faults (including child processes) since the start of the process. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

Total Size (Pages) The number of pages that the process has in real memory. Valid entry is an integer. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

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. Valid entry is an integer between -10,000 and 10,000.

VM Lock Pages (KB) The size (in kilobytes) of locked pages of the virtual memory. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

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. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

VM Size (KB) The size (in kilobytes) of the virtual memory. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

VM Size MB Virtual memory size in megabytes. This attribute displays as a floating point with a scale of 1. For example 5.2. Valid values can include Value_Exceeds_Maximum=2147483647 and Value_Exceeds_Minimum=-2147483648.

Process User Info Attributes

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

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

Effective Group ID The identifier of the effective group. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

Effective Group Name The effective group name. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

Effective Group Name (Unicode) The effective group name. Valid entry is a text string, with a maximum length of 64 bytes. This attribute is globalized (Unicode).

Effective User ID The identifier of the effective user. Valid entry is an integer.

Effective User Name The name of the effective user. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

Effective User Name (Unicode) The name of the effective user. Valid entry is a text string, with a maximum length of 64 bytes. This attribute is globalized (Unicode).

File System Group Name The name of the file system group. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

File System Group Name (Unicode) The name of the file system group. Valid entry is a text string, with a maximum length of 64 bytes. This attribute is globalized (Unicode).

File System Group ID The identifier of the file system group. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

File System User ID The identifier of the file system user. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

File System User Name The name of the file system user. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

File System User Name (Unicode) The name of the file system user. Valid entry is a text string, with a maximum length of 64 bytes. This attribute is globalized (Unicode).

Process Command Name (Unicode) The Process Command name (Unicode).

Process ID The identifier associated with the process. Valid entries: integers.

Process Parent ID The Parent Process ID.

Process State The state of the process (Sleeping, Disk, Running, Zombie, Trace, Dead, or N/A). Valid entry is an integer between -1 and 5, where:

0 = Sleeping

1 = Disk

2 = Running

3 = Zombie

4 = Trace

5 = Dead

-1 = Not_Available

Real Group ID The identifier of the real group. Valid entries: simple text string, alphanumeric with a maximum length 16 characters.

Real Group Name The name of the real group. Valid entries: simple text string, alphanumeric with a maximum length 16 characters.

Real Group Name (Unicode) The name of the real group. Valid entries: simple text string, with a maximum length 64 bytes. This attribute is globalized (Unicode).

Real User ID The identifier of the real user. Valid entry is an integer.

Real User Name The name of the real user. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

Real User Name (Unicode) The name of the real user. Valid entry is a text string, with a maximum length of 64 bytes. This attribute is globalized (Unicode).

Saved Group ID The identifier of the saved group. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

Saved Group Name The name of the saved group. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

Saved Group Name (Unicode) The name of the saved group. Valid entry is a text string, with a maximum length of 64 bytes.

Saved User ID The identifier of the saved user. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

Saved User Name The name of the saved user. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

Saved User Name (Unicode) The name of the saved user. Valid entry is a text string, with a maximum length of 64 bytes. This attribute is globalized (Unicode).

Session ID The session ID. 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:KLZ or deux.raleigh.ibm.com:KLZ.

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

VM Size (MB) Virtual Memory Size in Megabytes. This attribute displays as a floating point with a scale of 1. For example 5.2.

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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

RPC Client Calls The number of calls to the server made by the clients of the server. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

RPC Packets with Malformed Header The number of packets that were received at the server with header records that were not properly formatted. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

RPC Server Call Authorization Failures The number of packets that were received at the server with authorizations that were not valid. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

RPC Server Calls Rejected The number of calls made to the server, which were rejected. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

RPC Server Invalid Client Requests The number of packets that were received at the server, which had client requests that were not valid. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

RPC Total Server Calls Received The total number of calls made to the server (both valid and not valid). Valid entry is an integer. 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:KLZ or deux.raleigh.ibm.com:KLZ.

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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

Foreign Port The number of the foreign port. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Local Address The address of the local end of the socket, presented as a dotted ip address. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

Local Port The local port number. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Local Service Name The local port number translated to service name from /etc/services. Valid entry is an alphanumeric text string, with a maximum length of 64 characters.

Receive Queue Bytes The count of bytes not copied by the user program connected to this socket. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Send Queue Bytes The count of bytes not acknowledged by the remote host. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Socket Inode The inode used by the socket. Valid entry is an integer.

Socket Owner Name (Unicode) The user name associated with the user ID that owns or started the socket connection. Valid entry is a text string, with a maximum length of 64 bytes. This attribute is globalized (Unicode).

Socket Protocol Indicates the sockets using this protocol. "Total" includes UNIX® domain sockets not displayed here. Valid entry is an integer, where:

0 = TCP

1 = UDP

2 = RAW

3 = UNIX

-2 = Not_Available

Socket State The state of the socket. Valid entry is an integer, 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

12 = UNKNOWN

Socket UID The user ID of the owner of the socket. Valid entry is an integer.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Socket Protocol Indicates the sockets using this protocol. "Total" includes UNIX domain sockets not displayed here. Valid entry is an integer, where:

0 = TCP

1 = UDP

2 = RAW

3 = UNIX

4 = FRAG

-1 = TOTAL

-2 = NOT_AVAILABLE

Sockets in Use Sockets in use by protocol. Valid entry is an integer. 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:KLZ or deux.raleigh.ibm.com:KLZ.

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Swap Rate Attributes

The Swap Rate attributes feature swap space characteristics, including usage rates and days till data.

Days Until Swap Space Full The predicted number of days till swap space is completely used (moving average). Valid entry is an integer.

Low Water Mark for Free real memory (KB) The lowest level that Free real memory has reached, expressed in kilobytes. Valid entry is an integer. Note: -1 indicates N/A.

Minimum Days to Swap Full The minimum number of days till swap space is completely used (peak rate based). Valid entry is an integer.

Peak Swap Space Used (MB) The peak swap space used based on snap shots, expressed in megabytes. Valid entry is an integer.

Swap Space Used (MB) (Moving Average) The moving average of swap space used, expressed in megabytes. Valid entry is an integer.

Swap Space Used (bytes per hour) The swap space usage rate, expressed in bytes per hour. Valid entry is an integer.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Total Swap Space (MB) (Moving Average) The moving average of total swap space, expressed in megabytes. Valid entry is an integer.

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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

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: -1 indicates N/A.

Number of Processes in Zombie State Number of processes currently in Zombie State.

Number of User Logins The current number of users logged in. Valid entry is an integer. 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: -1 indicates N/A.

Pages Paged In Per Second The pages paged in per second. Calculated on a 30 second interval.

Pages Paged Out Per Second The pages paged out per second. Calculated on a 30 second interval.

Pages Swapped In The pages swapped in.

Pages Swapped In Per Second The pages swapped in per second. Calculated on a 30 second interval. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Pages Swapped Out The pages swapped out.

Pages Swapped Out Per Second The pages swapped out per second. Calculated on a 30 second interval. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Percent Change Context Switches Per Second The percent change in the number of context switches per second. Valid entry is an integer in the range -100 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

Percent Change Processes Created The percent change in the number of processes per second. Valid entry is an integer 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. Valid entry is an integer. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

System Load Last 1 Minute The load on the system for the last minute. Valid entry is an integer in the range 0 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

System Load Last 5 Minutes The load on the system for the last five minutes. Valid entry is an integer in the range 0 to 100. Note: -1 indicates Not_Available and -2 indicates Not_Collected.

System Load Last 15 Minutes The load on the system for the last fifteen minutes. Valid entry is an integer in the range 0 to 100. 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:KLZ or deux.raleigh.ibm.com:KLZ.

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. Valid entry is an integer.

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

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Total Number of Processes The total number of processes.

Total Pages Paged In The total pages paged in.

Total Pages Paged Out The total pages paged out.

User Login Attributes

The User Login attributes refer to user characteristics such as idle time, user name, location, and login time.

Hostname (From) The hostname associated with the login for the user. Valid entry is an alphanumeric text string, with a maximum length of 256 characters.

Idle Time The number of minutes that have passed since a user last entered a command. Valid entry is a numeric value 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. Valid entry is an alphanumeric text string, with a maximum length of 16 characters.

Login Time The date and time the user logged in. Valid entry is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Example: To express November 6, 2000, 1:05 p.m., enter 0981106130500000.

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

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

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

User Login PID The login ID of the user. Valid entry is an integer.

User Name The full name of a user. Valid entry is an alphanumeric text string, with a maximum length of 32 characters.

User Name (Unicode) The name of the user logging in to access the system. Valid entry is a text string up to 64 bytes. This attribute is globalized (Unicode).

VM Stats Attributes

The VM Stats attributes refer to memory characteristics such as the size of cached, free, and shared memory.

Available Virtual Storage (MB) The available virtual storage in MB. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Available Virtual Storage (Percent) The available virtual storage in percent. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Memory Cached (MB) The size (in megabytes) of physical memory cached. Valid entry is an integer. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Memory Free (MB) The size (in megabytes) of physical memory free. Valid entry is an integer. 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. Valid entry is an integer. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Memory Used (MB) The size (in megabytes) of physical memory used. Valid entry is an integer. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Real Memory Available (Percent) Available Real Memory in Percent. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Real Memory Used (Percent) Used Real Memory (Percent). Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Shared Memory (MB) The size (in megabytes) of physical memory shared. Valid entry is an integer. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Swap Space Available (Percent) Available Swap Space (Percent). Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Swap Space Free (MB) The size (in megabytes) of swap space free. Valid entry is an integer. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Swap Space Used (MB) The size (in megabytes) of swap space used. Valid entry is an integer. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

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:KLZ or deux.raleigh.ibm.com:KLZ.

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. The timestamp format for SCAN and STR functions is CYYMMDDHHMMSSmmm (as in 1020315064501000 for 03/15/02 06:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = millisecond

Total Memory (MB) The total size (in megabytes) of physical memory. Valid entry is an integer. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Total Swap Space (MB) The total size (in megabytes) of swap space. Valid entry is an integer. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Total Virtual Storage (MB) The total virtual storage (real plus swap storage) in MB. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Used Virtual Storage (MB) The used virtual storage in MB. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

Used Virtual Storage (Percent) The used virtual storage in percent. Note: the value -1 indicates Not Available and -2 indicates Not Collected.

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 upon the application environment that is being monitored. For example, if your attribute group is monitoring each processor on your machine and you have a dual processor machine, 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 8 on page 71 provides the following information required to calculate disk space for the Monitoring Agent for Linux OS:

- *DB table name* is the table name as it would appear in the warehouse database, if the attribute group is configured to be written to the warehouse.
- *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 8. Capacity planning for historical data logged by component linux

DB table name	Attribute group	Bytes per instance (agent)	Database bytes per instance (warehouse)	Aggregate bytes per instance (warehouse)
LNXALLUSR	Linux_All_Users	180	173	210
LNXCPU	Linux_CPU	184	249	665
LNXCPUAVG	Linux_CPU_Averages	208	346	1102
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	1660	1660	1697
LNXFILE	Linux_File_Information	3600	3636	3724
LNXFILPAT	Linux_File_Pattern	1660	1660	1697
LNXGROUP	Linux_Group	172	163	200
LNXPING	Linux_Host_Availability	252	255	343
LNXIOEXT	Linux_IO_Ext	276	474	1327
LNXMACHIN	Linux_Machine_Information	764	764	801
LNXNFS	Linux_NFS_Statistics	352	392	1740
LNXNET	Linux_Network	304	322	983
LNXOSCON	Linux_OS_Config	468	460	497
LNXPROC	Linux_Process	1168	1388	2805
LNXPUSR	Linux_Process_User_Info	1432	1469	1557
LNXRPC	Linux_RPC_Statistics	180	177	334
LNXSOCKD	Linux_Sockets_Detail	340	341	456
LNXSOCKS	Linux_Sockets_Status	160	152	228
LNXSWPRT	Linux_Swap_Rate	176	172	365
LNXSYS	Linux_System_Statistics	232	350	1194
LNXLOGIN	Linux_User_Login	552	552	589
LNXVM	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 6. Situations reference

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

About situations

A situation is a logical expression involving one or more system conditions. Situations are used to monitor the condition of systems in your network. You can manage situations from the Tivoli Enterprise Portal by using the Situation editor.

The IBM Tivoli Monitoring agents that you use to monitor your system environment are 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 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

More information about situations

The *IBM Tivoli Monitoring User's Guide* contains more information about predefined and custom situations and how to use them to respond to alerts.

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

Predefined situations

This monitoring agent contains the following predefined situations:

The remaining sections of this chapter contain descriptions of each of these predefined situations. The situations are organized alphabetically.

- Linux_Fragmented_File_System
- Linux_High_CPU_Overload
- Linux_High_CPU_System
- Linux_High_Packet_Collisions
- Linux_High_RPC_Retransmit
- Linux_High_Zombies
- Linux_Low_Pct_Inodes
- Linux_Low_percent_space
- Linux_Low_Space_Available
- Linux_Network_Status
- Linux_NFS_Buffer_High
- Linux_NFS_Getattr_High
- Linux_NFS_rdlink_high
- Linux_NFS_Read_High
- Linux_NFS_Writes_High
- Linux_Packets_Error
- Linux_Process_High_Cpu
- Linux_Process_stopped
- Linux_RPC_Bad_Calls
- Linux_System_Thrashing

Linux_Fragmented_File_System 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 Linux_Disk.Space_Used_Percent LT 85 AND
VALUE Linux_Disk.Inodes_Used_Percent GT 80
```

Linux_High_CPU_Overload 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 Linux_CPU.Idle_CPU LT 10.0 AND VALUE Linux_CPU.CPU_ID EQ Aggregate
```

Linux_High_CPU_System 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 Linux_CPU.CPU_ID EQ Aggregate AND VALUE
Linux_CPU.System_CPU GT 80.0
```

Linux_High_Packet_Collisions 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 Linux_Network.Collision_Percent GT 10
```

Linux_High_RPC_Retransmit 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 Linux_RPC_Statistics.RPC_Client_Calls_Retransmitted GT 10
```

Linux_High_Zombies 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 Linux_Process.State EQ Zombie AND COUNT Linux_Process.State GT 20
```

Linux_Low_Pct_Inodes 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 Linux_Disk.Inodes_Used_Percent GT 80
```

Linux_Low_percent_space 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 Linux_Disk.Space_Available_Percent LT 15
```

Linux_Low_Space_Available 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 Linux_Disk.Space_Available LT 7
```

Linux_Network_Status 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 Linux_Network.Interface_Status NE UP
```

Linux_NFS_Buffer_High 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 Linux_RPC_Statistics.RPC_Client_Calls_Retransmitted GT 60 AND
 PCTCHANGE Linux_RPC_Statistics.RPC_Client_Times_Authentication_Refreshed GT 5
```

Linux_NFS_Getattr_High 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 Linux_NFS_Statistics.NFS_Get_Attribute_Calls_Pct GT 40
```

Linux_NFS_rdlink_high 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 Linux_NFS_Statistics.NFS_Read_Link_Pct GT 10
```

Linux_NFS_Read_High 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 Linux_NFS_Statistics.NFS_Read_Calls_Pct GT 30
```

Linux_NFS_Writes_High 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 Linux_NFS_Statistics.NFS_Writes_Pct GT 15
```

Linux_Packets_Error 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 Linux_Network.Total_Error_Percent GT 10
```

Linux_Process_High_Cpu 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 Linux_Process.Busy_CPU_Pct GT 60.0
```

Linux_Process_stopped 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 Linux_Process.State NE Running AND
VALUE Linux_Process.State NE Sleeping
```

Linux_RPC_Bad_Calls 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 Linux_RPC_Statistics.RPC_Client_Calls_Retransmitted GT 30
) OR ( VALUE Linux_RPC_Statistics.RPC_Server_Calls_Rejected GT 30 ) )
```

Linux_System_Thrashing 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 Linux_System_Statistics.Pages_paged_out_per_sec GT 400.0 )
OR ( *VALUE Linux_System_Statistics.Pages_paged_in_per_sec GT 400.0 ) )
```

Chapter 7. Take Action commands reference

This chapter contains an overview of Take Action commands, references for detailed information about Take Action commands, and a description of the Take Actions command included in this monitoring agent.

About Take Action commands

Take Action commands can be run from the desktop or included in a situation or a policy.

When included in a situation, the command 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.

More information about Take Action commands

For more information about working with Take Action commands, see the *IBM Tivoli Monitoring User's Guide*.

Predefined Take Action commands

This monitoring agent contains the following Take Action command:

Sample_kill_Process

The remaining section of this chapter contains a description of this Take Action command. 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

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 8. Policies reference

This chapter contains an overview of policies and references for detailed information about policies.

About policies

Policies are an advanced automation technique for implementing more complex workflow strategies than you can create through simple automation.

A *policy* is a set of automated system processes that can perform actions, schedule work for users, or automate manual tasks. You use the Workflow Editor to design policies. You control the order in which the policy 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.

More information about policies

For more information about working with policies, see the *IBM Tivoli Monitoring User's Guide*.

For information about using the Workflow Editor, see the *IBM Tivoli Monitoring Administrator's Guide* or the Tivoli Enterprise Portal online help.

For a list of the policies for this monitoring agent and a description of each policy, refer to the "Predefined policies" section below and the information in that section for each individual policy.

Predefined policies

There are no predefined policies for this monitoring agent.

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.

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 9 contains a list of the time periods and the suffixes for the summary tables and views.

Table 9. Time periods and suffixes for summary tables and views

Data collection time period	Summary table suffixes	Summary view suffixes
Hourly	_H	_HV
Daily	_D	_DV
Weekly	_W	_WV
Monthly	_M	_MV

Table 9. Time periods and suffixes for summary tables and views (continued)

Data collection time period	Summary table suffixes	Summary view suffixes
Quarterly	_Q	_QV
Yearly	_Y	_YV

Table 10 shows the expansion to summary columns of some of the most commonly used attribute types.

Table 10. 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.

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 11.

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 12 on page 89. 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). 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 11. 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
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

Table 11. Overview of Distributed Monitoring migrated situations (continued)

Situation	IBM Tivoli Enterprise Console event class
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
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 table below 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 12. 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: • system_name: STRING • timestamp: STRING • user_name: STRING • login_pid: INTEGER • 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: • system_name: STRING • timestamp: STRING • disk_name: STRING • mount_point: STRING • size: INTEGER • space_used: INTEGER • space_available: INTEGER • total_inodes: INTEGER • inodes_used: INTEGER • inodes_free: INTEGER • space_used_percent: INTEGER • inodes_used_percent: INTEGER • fs_type: STRING • space_available_percent: STRING • mount_point_u: STRING • linux_vm_id: STRING • inodes_available_percent: INTEGER

Table 12. 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: • system_name: STRING • timestamp: INTEGER • disk_name: STRING • space_used: INTEGER • space_available: INTEGER • disk_usage_rate: INTEGER • highwater_du_rate: INTEGER • highwater_time: STRING • disk_usage_moving_average: INTEGER • 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 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Network	ITM_Linux_Network event class with these slots: • 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 • kbytes_received_count: INTEGER • bytes_received_per_sec: INTEGER • kbytes_transmitted_count: INTEGER • bytes_transmitted_per_sec: INTEGER • packets_received_count: INTEGER • packets_received_per_sec: INTEGER • input_errors: INTEGER • output_errors: INTEGER • packets_transmitted_per_sec: INTEGER • input_errors: INTEGER • output_errors: INTEGER • collisions: INTEGER • collision_rate: INTEGER • collision_percent: INTEGER • input_error_rate: INTEGER • output_error_rate: INTEGER • total_error_percent: INTEGER • input_packets_dropped: INTEGER • output_packets_dropped: INTEGER • input_fifo_buffer_overruns: INTEGER • output_fifo_buffer_overruns: INTEGER • packet_framing_errors: INTEGER • carrier_losses: INTEGER • linux_vm_id: STRING • input_error_percent: INTEGER • output_error_percent: INTEGER • device_type: INTEGER • device_type_enum: STRING

Table 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
User	ITM_Linux_CPU event class with these slots: • 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
Linux_CPU_Averages	ITM_Linux_CPU_Averages event class with these slots: • 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 • linux_vm_id: STRING

Table 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Process	ITM_Linux_Process event class with these slots: • 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 • resident_set_size: INTEGER • shared_memory: INTEGER • text_resident_size: INTEGER • shared_lib_set_size: INTEGER • data_set_size: INTEGER • dirty_pages: INTEGER • 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 • tot_minor_faults: INTEGER • tot_major_faults: INTEGER • 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

Table 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Process (continued)	• vm_size_mb: REAL • vm_lock_mb: REAL • vm_data_mb: REAL • vm_stack_mb: REAL • vm_exe_size_mb: REAL • vm_lib_size_mb: REAL • 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

Table 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_Process_User_Info	ITM_Linux_Process_User_Info event class with these slots: • system_name: STRING • timestamp: STRING • process_id: INTEGER • 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 • eff_user_id: INTEGER • saved_user_id: INTEGER • fs_user_id: INTEGER • real_group_id: INTEGER • eff_group_id: INTEGER • saved_group_id: INTEGER • file_sys_group_id: INTEGER • 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 • parent_process_id: INTEGER • state: INTEGER • state_enum: STRING • proc_cmd_line_u: STRING • process_command_name_u: STRING • vm_size_mb: REAL • terminal_device: STRING

Table 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_System_Statistics	ITM_Linux_System_Statistics event class with these slots: • system_name: STRING • timestamp: STRING • ctxt_switches_per_sec: INTEGER • ctxt_switches_per_sec_enum: STRING • pct_change_ctxt_switches: REAL • pct_change_ctxt_switches_enum: STRING • processes_per_sec: INTEGER • processes_per_sec_enum: STRING • 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 • linux_vm_id: STRING • pages_paged_in: INTEGER • pages_paged_in_per_sec: REAL • pages_paged_out: INTEGER • pages_paged_out_per_sec: REAL • pages_swapped_in: INTEGER • pages_swap_in_per_sec: REAL • pages_swapped_out: INTEGER • pages_swap_out_per_sec: REAL • 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 • processes_zombie_count: INTEGER

Table 12. 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: • system_name: STRING • timestamp: STRING • moving_total_swap_space: INTEGER • swap_space_used: INTEGER • swap_usage_rate: INTEGER • days_to_swap_space_full: INTEGER • peak_swap_space_used: INTEGER • days_to_peak_space_full: INTEGER • low_free_memory: INTEGER • low_free_memory_enum: STRING • linux_vm_id: STRING
Linux_VM_Stats	ITM_Linux_VM_Stats event class with these slots: • system_name: STRING • timestamp: STRING • total_swap_space: REAL • swap_space_used: REAL • swap_usage_free: REAL • total_memory: REAL • memory_used: REAL • memory_free: REAL • shared_memory: REAL • memory_in_buffers: REAL • memory_cached: REAL • linux_vm_id: STRING • total_virtual_storage: REAL • used_virtual_storage: REAL • available_virtual_storage: REAL • virtual_storage_pct_avail: INTEGER • virtual_storage_pct_used: INTEGER • real_memory_pct_used: INTEGER • real_memory_pct_avail: INTEGER • swap_pct_used: INTEGER • swap_pct_avail: INTEGER

Table 12. 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: • 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: • 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_inode: INTEGER • foreign_port: INTEGER • foreign_port_enum: STRING • socket_owner_name_u: STRING • linux_vm_id: STRING

Table 12. 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: • system_name: STRING • timestamp: STRING • transfers_per_sec: REAL • blk_rds_per_sec: REAL • blk_wrtn_per_sec: REAL • blks_read: INTEGER • blks_wrtn: INTEGER • dev_major: INTEGER • dev_minor: INTEGER • dev_name: STRING • linux_vm_id: STRING
Linux_IO_Ext	ITM_Linux_IO_Ext event class with these slots: • system_name: STRING • timestamp: STRING • device_name: STRING • read_reqm_per_sec: REAL • write_reqm_per_sec: REAL • read_req_per_sec: REAL • write_req_per_sec: REAL • read_sect_per_sec: REAL • write_sect_per_sec: REAL • avg_req_size: REAL • avg_req_queue_length: REAL • avg_wait_time: REAL • avg_svc_time: REAL • cpu_util: REAL • linux_vm_id: STRING • disk_read_percent: REAL • disk_write_percent: REAL • read_bytes_per_sec: REAL • write_bytes_per_sec: REAL • transfers_bytes_per_sec: REAL

Table 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_RPC_Statistics	ITM_Linux_RPC_Statistics event class with these slots: • 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 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_NFS_Statistics	ITM_Linux_NFS_Statistics event class with these slots: • 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 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_NFS_Statistics (continued)	• 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 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_NFS_Statistics (continued)	• 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	ITM_Linux_CPU_Config event class with these slots: • 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

Table 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_OS_Config	ITM_Linux_OS_Config event class with these slots: • system_name: STRING • timestamp: STRING • os_name: STRING • os_version: STRING • gcc_version: STRING • os_vendor: STRING
Linux_File_Information	ITM_Linux_File_Information event class with these slots: • system_name: STRING • timestamp: STRING • path_u: STRING • file_name_u: STRING • size_mb: REAL • 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 • modification_result: INTEGER • modification_result_enum: STRING
Linux_Host_Availability	ITM_Linux_Host_Availability event class with these slots: • system_name: STRING • timestamp: STRING • target_host: STRING • host_availability: INTEGER • host_availability_enum: STRING • response_time: REAL • response_time_enum: STRING

Table 12. Overview of attribute groups to event classes and slots (continued)

Attribute group	event classes and slots
Linux_File_Pattern	ITM_Linux_File_Pattern event class with these slots: • system_name: STRING • timestamp: STRING • file_name: STRING • match_pattern: STRING • match_option: INTEGER • match_option_enum: STRING • match_count: INTEGER • match_count_enum: STRING
Linux_File_Comparison	ITM_Linux_File_Comparison event class with these slots: • system_name: STRING • timestamp: STRING • file_name_1: STRING • file_name_2: STRING • file_compare_option: INTEGER • file_compare_option_enum: STRING • file_compare_result: INTEGER • file_compare_result_enum: STRING
Linux_All_Users	ITM_Linux_All_Users event class with these slots: • system_name: STRING • timestamp: STRING • name: STRING • user_id: INTEGER • password_null: INTEGER • password_null_enum: STRING • user_duplicated: INTEGER • user_duplicated_enum: STRING • user_sessions: INTEGER
Linux_Group	ITM_Linux_Group event class with these slots: • system_name: STRING • timestamp: STRING • group_name: STRING • group_id: INTEGER • group_duplicated: INTEGER • group_duplicated_enum: STRING

Table 12. 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: • 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

Appendix C. Problem determination

This appendix 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 appendix provides agent-specific problem determination information. See the *IBM Tivoli Monitoring Problem Determination Guide* for general problem determination information. Also see "Support for problem solving" on page 125 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 13. 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 108 for lists of all trace log files and their locations. See the *IBM Tivoli Monitoring User's Guide* for general information about the IBM Tivoli Monitoring environment.
Linux information	• Version number and patch level • Sample application data file (if monitoring a file)
Operating system	Operating system version number and patch level
Messages	Messages and other information displayed on the screen
Version numbers for IBM Tivoli Monitoring	Version number of the following members of the monitoring environment: • IBM Tivoli Monitoring. Also provide the patch level, if available. • IBM Tivoli Monitoring: 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 *install_dir*/bin directory, where *install_dir* is the directory path where you installed the monitoring agent.

Upload files for review to the following FTP site: `ftp.emea.ibm.com`. Log in as **anonymous** and place your files in the directory that corresponds to the IBM Tivoli Monitoring component that you use.

Built-in problem determination 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" 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 appendix 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 Problem Determination Guide* for general problem determination 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 109
- "Examples: using trace logs" on page 110
- "Setting RAS trace parameters" on page 111

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 14. Log file management on UNIX compared to log file management on Windows

Location of logs	Description
• 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	On Windows, the log file is overwritten each time the component starts. There is no automated method to archive previous RAS1 log files. **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.

Table 14. Log file management on UNIX compared to log file management on Windows (continued)

Location of logs	Description

• On a UNIX or Linux monitoring server
• On a UNIX or Linux computer where the monitoring agent is running
	On UNIX or Linux systems, because of the use of the &Timestamp; 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 time stamp as in these examples from a computer with a host name **f50pa2b**. The **1112097194** part of the name is the time stamp:
	`f50pa2b_lz_1112097194.log`
	`f50pa2b_lz_1112097194.pid60420`
	
	where *lz* is the unique, two-character code for Monitoring Agent for Linux OS.

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 15 can help you identify files that are relevant to your problem determination efforts.

Principal trace log files

Table 15 contains locations, file names, and descriptions of trace logs that can help determine the source of problems with agents.

Table 15. 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 *hostname*_lz_*instance*.log file is located in the *install_dir*/logs path.	Traces activity of the monitoring agent.
	The *.LG0 file is located in the following subdirectory of the *install_dir* path: /logs.	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 *previous* monitoring session: • Status of connectivity with the monitoring server. • Situations that were running. • The success or failure status of Take Action commands.
	The *take_action_name*.log file (where *take_action_name* is the name of the Take Action command) is located in the /logs subdirectory of the *install_dir* 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 start_command.log file.

Table 15. Trace log files for troubleshooting agents (continued)

System where log is located	File name and path	Description
On the Tivoli Enterprise Monitoring Server	The `candle_installation.log` file in the *`install_dir`*`/logs` path.	Provides details about products that are installed. **Note:** Trace logging is enabled by default. A configuration step is not required to enable this tracing.
	The `Warehouse_Configuration.log` file is located in the following path on Windows: *install_dir*\InstallITM.	Provides details about the configuration of data warehousing for historical reporting.
	The name of the RAS log file is as follows: • **On Windows:** *install_dir*\logs\ *`hostname`*`_ms_`*`timestamp`*`.log` • **On UNIX or Linux:** *`hostname`*`_ms_`*`timestamp`*`.log` and *`hostname`*`_ms_`*`timestamp`*`.pid`*`nnnnn`* in the *install_dir/logs* path, where *nnnnn* is the process ID number.	Traces activity on the monitoring server.
On the Tivoli Enterprise Portal Server	The name of the RAS log file is as follows: • **On Windows:** *install_dir*\logs\ *hostname*_cq_*timestamp*.log • **On UNIX or Linux:** *`hostname`*`_cq_`*`timestamp`*`.log` and *`hostname`*`_cq_`*`timestamp`*`.pid`*`nnnnn`* in the *install_dir/logs* path, where *nnnnn* is the process ID number.	Traces activity on the portal server.
	The `TEPS_ODBC.log` file is located in the following path on Windows: *`install_dir`*`\InstallITM`.	When you enable historical reporting, this log file traces the status of the warehouse proxy agent.
Definitions of variables: *timestamp* is time stamp whose format includes year (y), month (m), day (d), hour (h), and minute (m), as follows: **yyyymmdd hhmm** *install_dir* represents the directory path where you installed the IBM Tivoli Monitoring component. *install_dir* can represent a path on the computer that host the monitoring system, the monitoring agent, or the portal. *instance* refers to the name of the database instance that you are monitoring. *hostname* refers to the name of the computer on which the IBM Tivoli Monitoring component runs.		

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}kdcl0cl.c,105,"KDCL0_ClientLookup") status=1c020006,
    "location server unavailable", ncs/KDC1_STC_SERVER_UNAVAILABLE
(Thursday, August 11, 2005, 08:21:35-{94C}kraarreg.cpp,1157,"LookupProxy") Unable to connect to
    broker at ip.pipe:: status=0, "success", ncs/KDC1_STC_OK
(Thursday, August 11, 2005, 08:21:35-{94C}kraarreg.cpp,1402,"FindProxyUsingLocalLookup") Unable
    to find running CMS on CT_CMSLIST <IP.PIPE:#server1a>
```

Example two

The following excerpts from the trace log *for the monitoring server* show the status of an agent, identified here as "Remote node." The name of the computer where the agent is running is **SERVER5B**:

```
(42C039F9.0000-6A4:kpxreqhb.cpp,649,"HeartbeatInserter") Remote node SERVER5B: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 15 on page 109. 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:klz 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:klz ALL) (UNIT:kra ALL)'
     ```
  3. Restart the monitoring agent so that your changes take effect.

Problems and workarounds

The following sections provide symptoms and workarounds for problems that might occur with Monitoring Agent for Linux OS:

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 appendix provides agent-specific problem determination information. See the *IBM Tivoli Monitoring Problem Determination Guide* for general problem determination information.

Installation and configuration problem determination

This section provides tables that show solutions for installation, configuration, and uninstallation problems.

Agent upgrade and restart using non-root

The monitoring agent can run using a non-root user ID on UNIX and Linux systems. This can be done by running the **itmcmd agent start** command while logged in as a non-root user, and this can be done remotely by deploying the agent using the **Run As** option on the GUI or using the **_UNIX_STARTUP_.Username** option on the **tacmd addSystem** command line. If the agent is running using a non-root user ID, and then the agent is upgraded, restarted remotely, restarted as a result of a system reboot, or the **itmcmd agent start** is run using the root user ID, then the monitoring agent subsequently runs as the root user. To confirm the user ID that the monitoring agent is using, run the following command:

```
itm_install/bin/cinfo -r
```

If the agent is using root, and that is not the desired user ID, then use the following steps to restart the agent:

1. Log in as root.
2. Run the **itmcmd agent stop** command.
3. Log in (or 'su') to the user ID that you want the agent to run as.
4. Run the **itmcmd agent start** command.

If the agent was running as root because of a system reboot, then edit the startup file using the following steps so that the appropriate user ID is used the next time the system is rebooted:

1. Look at *install_dir*/registry/AutoStart, and get *NUM*.
2. Edit the autostart for your operating system:

   The location of the startup file is platform dependent as follows:
   - AIX®: /etc/rc.itm*NUM*
   - HP-UX: /sbin/init.d/ITMAgents*NUM*
   - Linux: /etc/init.d/ITMAgents*NUM*
   - Solaris: /etc/init.d/ITMAgents*NUM*
3. Add entries for your operating system using the following command:

   ```
   /usr/bin/su - instancename
   -c "install_dir/bin/itmcmd agent
   -h install_dir
   -o instancename
   start product_code"
   ```

   Where:

   instancename
   Name of the instance

   install_dir
   Name of the directory

   product_code
   2-character product code for the agent, for example, 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.

Table 16. 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.
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`*nnn* `unable to install agent`, where *nnn* is the release number.	Solve this problem as follows: 1. Run the following command before running any installation or configuration command for the agent: `export JAVA_COMPILER=NONE` 2. Install the following two RPM (Red Hat Package Manager) files: • **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. You can obtain the required RPM files from the CD for Red Hat As 4.0 s390x.
During a command-line installation, you choose to install a component that is already installed, and you see the following warning: `WARNING - you are about to install the SAME version of "`*component*`"` where *component* 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.	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: `install.sh warning: unarchive of "/cdrom/unix/cienv1.tar" may have failed`	This error is caused by low disk space. Although the install.sh script indicates that it is ready to install the agent software, the script considers the size of *all* tar files, not the size of all the files that are contained within the tar file.Run the **df -k** command to check whether the file systems have enough space to install agents.

Table 16. Problems and solutions for installation and configuration (continued)

Problem	Solution
Installing as root: The product has been installed as `root`, which is not recommended. Without re-installing the product, how can you change from `root` to a different user account?	When you install the product as `root` the files in the *`install_dir`* directory are owned by `root`. You must change the status of the files as follows: 1. While logged on as `root`, run the *`install_dir`*`/bin/UnSetRoot` script, as in this example: `UnSetRoot [ -h CANDLEHOME ] userID` The script resets all the files under the *`install_dir`* directory. 2. Run the *`install_dir`*`/bin/SetPerm` command. **SetPerm** sets `root` permission for specific IBM Tivoli Monitoring agent files. **About installing as root:** Normally, do not use the `root` user account to install or to start the Monitoring Agents for UNIX, for Linux, and for UNIX Logs. If you use the `root` user account to install the product, the files do not receive the correct permissions, and product behavior is unpredictable. To create a stable installation of the product, use one of the following options: • Create a user account with all the authority and permissions to install and run commands. For example, create a **tivoli** user account. —OR— • Use any user account other than `root` that has the required authority and permissions.
Cannot locate the **KDCB0_HOSTNAME** setting.	Go to *`install_dir`*`/config` 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: 1. Access the *`install_dir`*`/config/lz.ini` file, which is described in "Setting RAS trace parameters" on page 111. 2. Add the following line: `KBB_SIG1=trace –dumpoff`
Agents in the monitoring environment use different communication protocols. For example, some agents have security enabled and others do not.	Configure both the monitoring server and the Warehouse proxy server to accept multiple protocols, as described in the *IBM Tivoli Monitoring Installation and Setup Guide*.
Creating a firewall partition file: The partition file enables an agent to connect to the monitoring server through a firewall.	**How it works:** When the agents start, they search KDCPARTITION.TXT for the following matches: • An entry that matches the partition name **OUTSIDE**. • An entry that also includes a valid external address. For more information, see the *IBM Tivoli Monitoring Installation and Setup Guide*.
You see the following error: `Hub not registered with location broker. Error-code 1195.`	Confirm that the password within the Tivoli Enterprise Monitoring Server is correct.

Table 16. Problems and solutions for installation and configuration (continued)

Problem	Solution

The Monitoring Agent for Linux OS is started and running but not displaying data in the Tivoli Enterprise Portal.	Perform the following steps:
	1. Open the Manage Tivoli Enterprise Monitoring Services window.
	2. Right-click the name of the monitoring server.
	3. Select **Advanced > Add TEMS Application Support** in the pop-up menu. Add application support if any for any agent that is missing from the list. See in IBM Tivoli Monitoring Installation and Setup Guide 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:
	• Verify that the following entries are available in the *install_dir*\candle_installation.log file:*install_dir*\Install\IBM Tivoli Monitoring *timestamp*.log
	`... Browser Client support for ITM Agent for Linux`
	`... Desktop Client support for ITM Agent for Linux`
	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 IBM Tivoli Monitoring Installation and Setup Guide for more information on adding application support.
	• Verify that the following files are available in the directory:
	`install_dir\ATTRLIB\klz.atr`
	`install_dir\CNPS\CMSATR\klz.atr`
	`install_dir\SQLLIB\klz.sql`
	`install_dir\CNPS\SQLLIB\klz.sql`

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`.	Install the agent's application support files on the Tivoli Enterprise Monitoring Server, using the following steps: 1. Open the Manage Tivoli Enterprise Monitoring Services window. 2. Right-click the name of the monitoring server. 3. Select **Advanced > Add TEMS Application Support** in the pop-up menu. Add application support if any for any agent that is missing from the list. See in IBM Tivoli Monitoring Installation and Setup Guide for more information on adding application support. Ensure that the agent's application support files are pushed to the system that houses the Warehouse Proxy Agent. The Warehouse Proxy must be able to access the short attribute names for tables and columns. That way, if the longer versions of these names exceed the limits of the Warehouse database, the shorter names can be substituted.

Table 16. Problems and solutions for installation and configuration (continued)

Problem	Solution
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 17. General problems and solutions for uninstallation

Problem	Solution
The way to remove inactive managed systems (systems whose status is OFFLINE) from the Enterprise navigation tree in the portal is not obvious.	When you want to remove a managed system from the navigation tree, complete the following steps: 1. Click **Enterprise** in the navigation tree. 2. Right-click **Workspace -> Managed System Status**. 3. Right-click the offline managed system and select **Clear offline entry**.

Unique names for monitoring components

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`.

   Note: When you modify the **lz.ini** file, your configuration changes affect only the instance Monitoring Agent for Linux OS that is running on the computer. If you want your configuration changes to affect all agents that run on the computer, modify the *install_dir*`/config/env.config` file.
2. Find the line the begins with **CTIRA_HOSTNAME=**.
3. Type a new name for host name that is a unique, shorter name for the host computer. The final concatenated name including the subsystem name, new host name, and 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 problem determination

This section lists problems that might occur with agents.

This appendix provides agent-specific problem determination information. See the *IBM Tivoli Monitoring Problem Determination Guide* for general problem determination information.

Table 18. Agent problems and solutions

Problem	Solution
A configured and running instance of the monitoring agent is not displayed in the Tivoli Enterprise Portal, but other instances of the monitoring agent on the same system do appear in the portal.	Tivoli Monitoring products use Remote Procedure Call (RPC) to define and control product behavior. RPC is the mechanism that allows a client process to make a subroutine call (such as GetTimeOfDay or ShutdownServer) to a server process somewhere in the network. Tivoli processes can be configured to use TCP/UDP, TCP/IP, SNA, and SSL as the desired protocol (or delivery mechanism) for RPCs. "IP.PIPE" is the name given to Tivoli TCP/IP protocol for RPCs. The RPCs are socket-based operations that use TCP/IP ports to form socket addresses. IP.PIPE implements virtual sockets and multiplexes all virtual socket traffic across a single physical TCP/IP port (visible from the netstat command). A Tivoli process derives the physical port for IP.PIPE communications based on the configured, well-known port for the HUB Tivoli Enterprise Monitoring Server. (This well-known port or BASE_PORT is configured using the 'PORT:' keyword on the KDC_FAMILIES / KDE_TRANSPORT environment variable and defaults to '1918'.) The physical port allocation method is defined as (BASE_PORT + 4096*N) where N=0 for a Tivoli Enterprise Monitoring Server process and N={1, 2, ..., 15} for a non-Tivoli Enterprise Monitoring Server. Two architectural limits result as a consequence of the physical port allocation method: • No more than one Tivoli Enterprise Monitoring Server reporting to a specific Tivoli Enterprise Monitoring Server HUB can be active on a system image. • No more that 15 IP.PIPE processes can be active on a single system image. A single system image can support any number of Tivoli Enterprise Monitoring Server processes (address spaces) provided that each Tivoli Enterprise Monitoring Server on that image reports to a different HUB. By definition, there is one Tivoli Enterprise Monitoring Server HUB per monitoring Enterprise, so this architecture limit has been simplified to one Tivoli Enterprise Monitoring Server per system image. No more that 15 IP.PIPE processes or address spaces can be active on a single system image. With the first limit expressed above, this second limitation refers specifically to Tivoli Enterprise Monitoring Agent processes: no more that 15 agents per system image. This limitation can be circumvented (at current maintenance levels, IBM Tivoli Monitoring V6.1 Fix Pack 4 and later) if the Tivoli Enterprise Monitoring Agent process is configured to use EPHEMERAL IP.PIPE. (This is IP.PIPE configured with the 'EPHEMERAL:Y' keyword in the KDC_FAMILIES / KDE_TRANSPORT environment variable). There is no limitation to the number of ephemeral IP.PIPE connections per system image. However, EPHEMERAL endpoints are restricted: data warehousing cannot be performed on an ephemeral endpoint.
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.	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 *Configuring Tivoli Enterprise Monitoring Server on z/OS®*.

Table 18. Agent problems and solutions (continued)

Problem	Solution
The agent's process, **klzagent** uses a large amount of system resources.	In most cases, the problem occurs during the backup. Any one of the following scenarios can cause this problem. **The agent is running during the backup** After backing up, the agent is started during system startup. **Multiple agents are running at the same time.** The computer that hosts the Tivoli Enterprise Monitoring Server was rebooted and the agent has been installed by the `root` user account. **The agent is running during the backup** During the backup, some of the service might be interrupted or not be available or locked for some amount of time. While the backup process is going on, the Monitoring Agent for 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. **The agent is started during system boot up:** If you use scripts to stop and start the agent, do not start the agent from an **init** process script when you restart the system. The computer that hosts the Tivoli Enterprise Monitoring Server was rebooted and the agent has been installed by the `root` user account. Verify whether the log file has the following information: `Unable to find running Tivoli Enterprise Monitoring Server on CMSLIST`
Attributes do not allow non-ASCII input in the situation editor.	None. Any attribute that does not include "(Unicode)" -- for example, "Description (Unicode)" might support only ASCII characters.
In the User workspace, data does not show up in the User Login Information (table view).	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.

Tivoli Enterprise Portal problem determination

Table 19 lists problems that might occur with the Tivoli Enterprise Portal. This appendix provides agent-specific problem determination information. See the *IBM Tivoli Monitoring Problem Determination Guide* for general problem determination information.

Table 19. Tivoli Enterprise Portal problems and solutions

Problem	Solution
Historical data collection is unavailable because of incorrect queries in the Tivoli Enterprise Portal.	The column, Sort By, Group By, and First/Last functions are not compatible with the historical data collection feature. Use of these advanced functions will make a query ineligible for historical data collection. Even if data collection has been started, you cannot use the time span feature if the query for the chart or table includes any column functions or advanced query options (Sort By, Group By, First / Last). To ensure support of historical data collection, do not use the Sort By, Group By, or First/Last functions in your queries. See the *IBM Tivoli Monitoring Administrator's Guide* the Tivoli Enterprise Portal online Help for information on the Historical Data Collection function.

Table 19. Tivoli Enterprise Portal problems and solutions (continued)

Problem	Solution
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.

Problem determination for remote deployment

Table 20 lists problems that might occur with remote deployment. This appendix provides agent-specific problem determination information. See the *IBM Tivoli Monitoring Problem Determination Guide* for general problem determination information.

This section describes problems and solutions for remote deployment and removal of agent software Agent Remote Deploy:

Table 20. 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.

Situation problem determination

This section provides information about both general situation problems and problems with the configuration of situations. See the *IBM Tivoli Monitoring Problem Determination Guide* for more information about problem determination for situations.

General situation problems

Table 21 lists problems that might occur with specific situations.

Table 21. Specific situation problems and solutions

Problem	Solution
You want to change the appearance of situations when they are displayed in a Workspace view.	1. Right-click an item in the Navigation tree. 2. Select **Situations** in the pop-up menu. The Situation Editor window is displayed. 3. Select the situation that you want to modify. 4. Use the **Status** pull-down menu in the lower right of the window to set the status and appearance of the Situation when it triggers. **Note:** This status setting is not related to severity settings in IBM Tivoli Enterprise Console.

Table 21. Specific situation problems and solutions (continued)

Problem	Solution
Situations are triggered in the Tivoli Enterprise Monitoring Server, but events for the situation are not sent to the Tivoli Enterprise Console server. The Tivoli Enterprise Monitoring Server is properly configured for event forwarding, and events for many other situations are sent to the event server.	This condition can occur when a situation is only monitoring the status of other situations. The event forwarding function requires an attribute group reference in the situation in order to determine the correct event class to use in the event. When the situation only monitors other situations, no attribute groups are defined and the event class cannot be determined. Because the event class cannot be determined, no event is sent. This is a limitation of the Tivoli Enterprise Monitoring Server event forwarding function. Situations that only monitor other situations do not send events to the event server.
Monitoring activity requires too much disk space.	Check the RAS trace logging settings that are described in "Setting RAS trace parameters" on page 111. For example, trace logs grow rapidly when you apply the **ALL** logging option.
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 #'Linux_VM_Stats.Total_Memory' / 10`	This formula is incorrect because situation predicates support only logical operators. Your formulas cannot have mathematical operators. **Note:** The Situation Editor provides alternatives to math operators. Regarding the example, you can select **% Memory Free** attribute and avoid the need for math operators.
If you are running a Version 350 Monitoring Agent for 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.	To enable Unicode and other features, upgrade the monitoring agent to IBM Tivoli Monitoring, Version 6.1.0.
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.	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: • 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
You see the 'Unable to get attribute name' error in the Tivoli Enterprise Monitoring Server log after creating a situation.	Install the agent's application support files on the Tivoli Enterprise Monitoring Server, using the following steps: 1. Open the Manage Tivoli Enterprise Monitoring Services window. 2. Right-click the name of the monitoring server. 3. Select **Advanced > Add TEMS Application Support** in the pop-up menu. Add application support if any for any agent that is missing from the list. See in IBM Tivoli Monitoring Installation and Setup Guide for more information on adding application support.
Events received at the Tivoli Enterprise Console server from IBM Tivoli Monitoring do not have values for all event attributes (slots) even though the values are visible in workspace views.	The problem is due to a limitation in the IBM Tivoli Monitoring interface code that generates Tivoli Enterprise Console events from situations. The situation results are provided in a chain of buffers of 3000 bytes each. The interface code currently extracts event information from only the first buffer. When situations or agent table data expands into a second buffer, this additional data is not examined, and it is not included in events sent to the Tivoli Enterprise Console server.

Table 21. 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: 1. Ensure that you have the IBM Tivoli Monitoring 6.2 Event Sync installed on your Tivoli Enterprise Console server. 2. Obtain updated baroc files from IBM Tivoli Monitoring 6.2 for the monitoring agent's events. Updated baroc files are on the Tivoli Enterprise Monitoring Server in the *CANDLEHOME*/CMS/TECLIB/itm5migr directory.
You are receiving Tivoli Business Systems Management events that cannot be associated due to application_oid and application_class not being set.	The problem is due to IBM Tivoli Monitoring 6.2 sending Tivoli Enterprise Console events for IBM Tivoli Monitoring 5.x migrated situations. These events are not able to set the cited slot values. Replace the *agent_name*_forward_tbsm_event_cb.sh script on the Tivoli Enterprise Console server with the version of this file from the Tivoli Enterprise Monitoring Server in the *CANDLEHOME*/CMS/TECLIB/itm5migr directory.

Problems with configuration of situations

Table 22 lists problems that might occur with situations.

This section provides information for problem determination for agents. Be sure to consult the *IBM Tivoli Monitoring Problem Determination Guide* for more general problem determination information.

Table 22. Problems with configuring situations that you solve in the Situation Editor

Problem	Solution
Note: To get started with the solutions in this section, perform these steps: 1. Launch the Tivoli Enterprise Portal. 2. Click **Edit > Situation Editor**. 3. In the tree view, choose the agent whose situation you want to modify. 4. Choose the situation in the list. The Situation Editor view is displayed.	
The situation for a specific agent is not visible in the Tivoli Enterprise Portal.	Open the Situation Editor. Access the All managed servers view. If the situation is absent, confirm that application support for Monitoring Agent for Linux OS has been added to the monitoring server. If not, add application support to the server, as described in the *IBM Tivoli Monitoring Installation and Setup Guide.*
The monitoring interval is too long.	Access the Situation Editor view for the situation that you want to modify. Check the **Sampling interval** area in the **Formula** tab. Adjust the time interval as needed.
The situation did not activate at startup.	Manually recycle the situation as follows: 1. Right-click the situation and choose **Stop Situation**. 2. Right-click the situation and choose **Start Situation**. **Note:** You can permanently avoid this problem by placing a check mark in the **Run at Startup** option of the Situation Editor view for a specific situation.
The situation is not displayed.	Click the **Action** tab and check whether the situation has an automated corrective action. This action can occur directly or through a policy. The situation might be resolving so quickly that you do not see the event or the update in the graphical user interface.
An Alert event has not occurred even though the predicate has been properly specified.	Check the logs, reports, and workspaces.
A situation fires on an unexpected managed object.	Confirm that you have distributed and started the situation on the correct managed system.

Table 22. Problems with configuring situations that you solve in the Situation Editor (continued)

Problem	Solution
The product did not distribute the situation to a managed system.	Click the **Distribution** tab and check the distribution settings for the situation.
The situation does not fire. Incorrect predicates are present in the formula that defines the situation. For example, the managed object shows a state that normally triggers a monitoring event, but the situation is not true because the wrong attribute is specified in the formula.	In the **Formula** tab, analyze predicates as follows: 1. Click the *fx* icon in the upper-right corner of the Formula area. The Show formula window is displayed. a. Confirm the following details in the **Formula** area at the top of the window: • The attributes that you intend to monitor are specified in the formula. • The situations that you intend to monitor are specified in the formula. • The logical operators in the formula match your monitoring goal. • The numerical values in the formula match your monitoring goal. b. (*Optional*) Click the **Show detailed formula** check box in the lower left of the window to see the original names of attributes in the application or operating system that you are monitoring. c. Click **OK** to dismiss the Show formula window. 2. (*Optional*) In the Formula area of the **Formula** tab, temporarily assign numerical values that will immediately trigger a monitoring event. The triggering of the event confirms that other predicates in the formula are valid. **Note:** After you complete this test, you must restore the numerical values to valid levels so that you do not generate excessive monitoring data based on your temporary settings.

Table 23. Problems with configuration of situations that you solve in the Workspace area

Problem	Solution
Situation events are not displayed in the Events Console view of the workspace.	Associate the situation with a workspace. **Note:** The situation does not need to be displayed in the workspace. It is sufficient that the situation be associated with any workspace.
You do not have access to a situation.	**Note:** You must have administrator privileges to perform these steps. 1. Select **Edit** > **Administer Users** to access the Administer Users window. 2. In the Users area, select the user whose privileges you want to modify. 3. In the Permissions tab, Applications tab, and Navigator Views tab, select the permissions or privileges that correspond to the user's role. 4. Click **OK**.
A managed system seems to be offline.	1. Select Physical View and highlight the Enterprise Level of the navigator tree. 2. Select **View** > **Workspace** > **Managed System Status** to see a list of managed systems and their status. 3. If a system is offline, check network connectivity and status of the specific system or application.

Table 24. 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.

Table 24. Problems with configuration of situations that you solve in the Manage Tivoli Enterprise Monitoring Services window (continued)

Problem	Solution
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.

Support for problem solving

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

- "Using IBM Support Assistant"
- "Obtaining fixes"
- "Contacting IBM Software Support" on page 126

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 Version 3, 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 IBM Tivoli Monitoring:

1. Start the IBM Support Assistant application.
2. Select **Updater** on the Welcome page.
3. Select **New Properties and Tools**.
4. Under Tivoli, select **IBM Tivoli Monitoring 6.2**, and then click **Install**. Be sure to read the license and description.
5. Restart the IBM Support Assistant.

Obtaining fixes

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

1. Go to the IBM Software Support Web site at http://www.ibm.com/software/support.
2. Under **Select a brand and/or product**, select **Tivoli** and click **Go**.

3. Under **Select a category**, select a product and click **Go**.
4. Under **Download**, click the name of a fix to read its description and, optionally, to download it.

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

Contacting IBM Software Support

IBM Software Support provides assistance with product defects. The easiest way to obtain that assistance is to open a PMR or ETR directly from the IBM Support Assistant (see "Using IBM Support Assistant" on page 125).

Before contacting IBM Software Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenance contract that you need depends on the type of product you have:

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus®, and Rational® products, as well as DB2 and WebSphere® products that run on Windows or UNIX operating systems), enroll in Passport Advantage® in one of the following ways:

  Online

  Go to the Passport Advantage Web site at http://www-306.ibm.com/software/howtobuy/passportadvantage/pao_customers.htm .

  By phone

  For the phone number to call in your country, go to the IBM Software Support Web site at http://techsupport.services.ibm.com/guides/contacts.html and click the name of your geographic region.
- For customers with Subscription and Support (S & S) contracts, go to the Software Service Request Web site at https://techsupport.services.ibm.com/ssr/login.
- For customers with IBMLink™, CATIA, Linux, OS/390®, iSeries®, pSeries®, zSeries®, and other support agreements, go to the IBM Support Line Web site at http://www.ibm.com/services/us/index.wss/so/its/a1000030/dt006.
- For IBM eServer™ software products (including, but not limited to, DB2 and WebSphere products that run in zSeries, pSeries, and iSeries environments), you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for eServer software products, go to the IBM Technical Support Advantage Web site at http://www.ibm.com/servers/eserver/techsupport.html.

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. From other countries, go to the contacts page of the *IBM Software Support Handbook* on the Web at http://techsupport.services.ibm.com/guides/contacts.html and click the name of your geographic region for phone numbers of people who provide support for your location.

To contact IBM Software support, follow these steps:

1. "Determining the business impact" on page 127
2. "Describing problems and gathering information" on page 127
3. "Submitting problems" on page 127

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Use the following criteria TO understand and assess the business impact of the problem that you are reporting:

Severity 1
The problem has a *critical* business impact. You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.

Severity 2
The problem has a *significant* business impact. The program is usable, but it is severely limited.

Severity 3
The problem has *some* business impact. The program is usable, but less significant features (not critical to operations) are unavailable.

Severity 4
The problem has *minimal* business impact. The problem causes little impact on operations, or a reasonable circumvention to the problem was implemented.

Describing problems and gathering information

When describing a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently. To save time, know the answers to these questions:

- Which software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can you re-create the problem? If so, what steps were performed to re-create the problem?
- Did you make any changes to the system? For example, did you make changes to the hardware, operating system, networking software, and so on.
- Are you currently using a workaround for the problem? If so, be prepared to explain the workaround when you report the problem.

Submitting problems

You can submit your problem to IBM Software Support in one of two ways:

Online
Click **Submit and track problems** on the IBM Software Support site athttp://www.ibm.com/software/support/probsub.html. Type your information into the appropriate problem submission form.

By phone
For the phone number to call in your country, go to the contacts page of the *IBM Software Support Handbook* at http://techsupport.services.ibm.com/guides/contacts.html and click the name of your geographic region.

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the Software Support Web site daily, so that other users who experience the same problem can benefit from the same resolution.

Appendix D. Documentation library

This appendix contains information about the publications related to the Monitoring Agent for Linux OS. These publications are listed in the following categories:

- Monitoring Agent for Linux OS library
- Prerequisite publications
- Related publications

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

To find a list of new and changed publications, click **What's new** on the Welcome page of the IBM Tivoli Monitoring and OMEGAMON XE Information Center. To find publications from the previous version of a product, click **Previous information centers** on the Welcome page for the product.

Monitoring Agent for Linux OS library

There is one document specific to the Monitoring Agent for Linux OS: *IBM Tivoli Monitoring: Linux OS Agent User's Guide*. This user's guide provides agent-specific reference and problem determination information for configuring and using the IBM Tivoli Monitoring for Linux OS Agent.

Use the configuration chapter in this guide with the *IBM Tivoli Monitoring Installation and Setup Guide* to set up the software.

Use the information in this guide with the *IBM Tivoli Monitoring User's Guide* to monitor Linux resources.

Prerequisite publications

To use the information in this publication effectively, you must have some prerequisite knowledge, which you can obtain from the following IBM Tivoli Monitoring publications:

- *Exploring IBM Tivoli Monitoring*
- *IBM Tivoli Monitoring Administrator's Guide*
- *IBM Tivoli Monitoring Agent Builder User's Guide*
- *IBM Tivoli Monitoring Command Reference*
- *IBM Tivoli Monitoring Installation and Setup Guide*
- *IBM Tivoli Monitoring: Messages*
- *IBM Tivoli Monitoring Migration Toolkit User's Guide*
- *IBM Tivoli Monitoring Problem Determination Guide*
- *IBM Tivoli Monitoring: Upgrading from Tivoli Distributed Monitoring*
- *IBM Tivoli Monitoring User's Guide*
- *IBM Tivoli Monitoring: Upgrading from V5.1.2*

- *IBM Tivoli Monitoring Configuring Tivoli Enterprise Monitoring Server on z/OS*
- *IBM Tivoli Monitoring: Windows OS Agent User's Guide*
- *IBM Tivoli Monitoring: UNIX OS Agent User's Guide*
- *IBM Tivoli Monitoring: Linux OS Agent User's Guide*
- *IBM Tivoli Monitoring: i5/OS Agent User's Guide*
- *IBM Tivoli Monitoring: UNIX Log Agent User's Guide*
- *IBM Tivoli Monitoring Universal Agent User's Guide*
- *IBM Tivoli Monitoring Universal Agent API and Command Programming Reference Guide*
- *Introducing IBM Tivoli Monitoring Version 6.1*

Related publications

The following documents also provide useful information:

- *IBM Tivoli Enterprise Console Adapters Guide*
- *IBM Tivoli Enterprise Console Event Integration Facility User's Guide*
- *IBM Tivoli Enterprise Console Reference Manual*
- *IBM Tivoli Enterprise Console Rule Builder's Guide*

Other sources of documentation

You can also obtain technical documentation about Tivoli Monitoring and OMEGAMON XE products from the following sources:

- IBM Tivoli Open Process Automation Library (OPAL)

  http://www.ibm.com/software/tivoli/opal

  OPAL is an online catalog that contains integration documentation as well as other downloadable product extensions. This library is updated daily.
- Redbooks

  http://www.redbooks.ibm.com/

  IBM Redbooks®, Redpapers, and Redbooks Technotes provide information about products from platform and solution perspectives.
- Technotes

  You can find Technotes through the IBM Software Support Web site at http://www.ibm.com/software/support/probsub.html, or more directly through your product Web site, which contains a link to Technotes (under **Solve a problem**).

  Technotes provide the latest information about known product limitations and workarounds.

Appendix E. Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in this product enable users to do the following:

- Use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. Consult the product documentation of the assistive technology for details on using those technologies with this product.
- Operate specific or equivalent features using only the keyboard.
- Magnify what is displayed on the screen.

In addition, the product documentation was modified to include the following features to aid accessibility:

- All documentation is available in both HTML and convertible PDF formats to give the maximum opportunity for users to apply screen-reader software.
- All images in the documentation are provided with alternative text so that users with vision impairments can understand the contents of the images.

Navigating the interface using the keyboard

Standard shortcut and accelerator keys are used by the product and are documented by the operating system. Refer to the documentation provided by your operating system for more information.

Magnifying what is displayed on the screen

You can enlarge information on the product windows using facilities provided by the operating systems on which the product is run. For example, in a Microsoft Windows environment, you can lower the resolution of the screen to enlarge the font sizes of the text on the screen. Refer to the documentation provided by your operating system for more information.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 systems. 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 system 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.

If you are viewing this information in softcopy form, the photographs and color illustrations might not appear.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

Trademarks

IBM, the IBM logo, IBMLink, AIX, Candle, DB2, developerWorks®, eServer, i5/OS, iSeries, Lotus, OMEGAMON, OS/390, OS/400®, Passport Advantage, pSeries, Rational, Redbooks, Tivoli, the Tivoli logo, Tivoli Enterprise, Tivoli Enterprise Console, VTAM®, WebSphere, z/OS, z/VM®, and zSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft, Windows, and Windows NT® are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside®, Intel Inside logo, Intel Centrino®, Intel Centrino logo, Celeron®, Intel Xeon®, Intel SpeedStep®, Itanium, and Pentium® are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

Index

L

M

N

O

P

Q

R

S

T

U

V

W

Printed in USA

SC32-9447-01