

IBM Tivoli Monitoring for Virtual Environments Agent for
Linux Kernel-based Virtual Machines
Version 7.2

User's Guide



IBM Tivoli Monitoring for Virtual Environments Agent for
Linux Kernel-based Virtual Machines
Version 7.2

User's Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 99.

This edition applies to version 7.2 of IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines (product number 5724-L92) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2010, 2013.**

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

Contents

Figures	v	Situation descriptions	54
Tables	vii	Linux Kernel-based Virtual Machines Navigator item	54
Chapter 1. Overview of the agent	1	Host Navigator item	55
IBM Tivoli Monitoring	1	Storage Pool Navigator item.	57
Functions of the monitoring agent	1	Virtual Machine Navigator item	58
New in this release	1	Chapter 6. Take Action commands reference.	61
Components of the IBM Tivoli Monitoring environment	2	Predefined Take Action commands	61
Agent Management Services	3	Chapter 7. Policies reference	63
User interface options	3	Predefined policies	63
Chapter 2. Agent installation and configuration	5	Chapter 8. Troubleshooting	65
Requirements	5	Trace logging	66
Language pack installation	5	Overview of log file management	66
Installing language packs on Windows systems.	5	Principal trace log files	67
Installing language packs on UNIX or Linux systems	6	Examples: Using trace logs	70
Silent installation of language packs on Windows, UNIX, or Linux systems	6	RAS trace parameters	71
Agent-specific installation and configuration.	8	Dynamic modification of trace settings	73
Virtualization hosts	8	Setting trace parameters for the Tivoli Enterprise Console server	76
Configuration values	11	Problems and workarounds	77
Remote installation and configuration	14	Installation and configuration troubleshooting.	77
Chapter 3. Workspaces reference	15	Remote deployment troubleshooting	80
Predefined workspaces	15	Agent troubleshooting.	81
Workspace descriptions	16	Workspace troubleshooting	83
Linux Kernel-based Virtual Machines Navigator item	16	Situation troubleshooting.	86
Host Navigator item	17	Take Action commands troubleshooting	88
Storage Pool Navigator item.	18	Support information	88
Virtual Machine Navigator item	19	Informational, warning, and error messages overview	89
Chapter 4. Attributes reference	21	Message format	89
Attribute groups for the monitoring agent	21	Agent messages	90
Attributes in each attribute group	22	Appendix A. Event mapping	91
Host CPU attribute group	22	Appendix B. Documentation library	97
Host Memory attribute group	27	Prerequisite publications	97
Hosts attribute group	31	Related publications	98
Performance Object Status attribute group	35	Other sources of documentation	98
Scheduler Parameters attribute group.	40	Notices	99
Storage Pools attribute group	42	Trademarks	101
Virtual Machines attribute group	44	Index	103
Disk capacity planning for historical data	51		
Chapter 5. Situations reference	53		
Predefined situations	53		

Figures

1. Agent Configuration: Data Provider 11
2. Agent Configuration: Hypervisor 13

Tables

1. Capacity planning for historical data logged by the Linux Kernel-based Virtual Machines agent	51	6. Remote deployment problems and solutions	80
2. Information to gather before contacting IBM Software Support	65	7. Agent problems and solutions	81
3. Trace log files for troubleshooting agents	67	8. Workspace problems and solutions.	84
4. Problems and solutions for installation and configuration	77	9. Situation problems and solutions	86
5. General problems and solutions for uninstallation	78	10. Take Action commands problems and solutions	88

Chapter 1. Overview of the agent

The IBM® Tivoli® Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines (product code V1) provides you with the capability to monitor Linux Kernel-based Virtual Machines.

IBM Tivoli Monitoring is the base software for the Linux Kernel-based Virtual Machines agent.

IBM Tivoli Monitoring

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 achieve the following tasks:

- 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 take actions, schedule work, and automate manual tasks.

The Tivoli Enterprise Portal is the interface for IBM Tivoli Monitoring products. You can use the consolidated view of your environment as seen in the Tivoli Enterprise Portal to monitor and resolve performance issues throughout the enterprise.

See the IBM Tivoli Monitoring publications listed in “Prerequisite publications” on page 97 for complete information about IBM Tivoli Monitoring and the Tivoli Enterprise Portal.

Functions of the monitoring agent

Display health and performance of Linux KVM host hypervisor systems and guest virtual machines

You can use the Linux Kernel-based Virtual Machines agent to visualize host capacity and virtual machine consumption in Linux KVM systems.

New in this release

For version 7.2 of the Linux Kernel-based Virtual Machines agent, the following enhancements have been made since version 7.1:

- New attribute groups
 - KV1_SCHEDULER_PARAMETERS
For more information, see “Scheduler Parameters attribute group” on page 40
- New or changed attributes in the following attribute groups
 - KV1_VIRTUAL_MACHINES
CPU Shares attribute added. For more information, see “Virtual Machines attribute group” on page 44
- Introduction of Tivoli Common Reporting.
Introduction of pre-defined Cognos® reports for the Linux Kernel-based Virtual Machines agent. For more information, see the *IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning Version 7.2 User’s Guide*.

- Changes related to system requirements. See the information about system requirements in Software product compatibility reports (<http://publib.boulder.ibm.com/infocenter/prodguid/v1r0/clarity/index.html>).
- Updated kv1.baroc file to support IBM Tivoli Enterprise Console® event mapping changes.

Components of the IBM Tivoli Monitoring environment

After you install and set up the Linux Kernel-based Virtual Machines agent, you have an environment that contains the client, server, and monitoring agent implementation for Tivoli Monitoring.

This Tivoli Monitoring environment contains the following components:

Tivoli Enterprise Portal client

The portal has a user interface based on Java™ for viewing and monitoring your enterprise.

Tivoli Enterprise Portal Server

The portal server is placed between the client and the Tivoli Enterprise Monitoring Server and enables retrieval, manipulation, and analysis of data from the monitoring agents. The Tivoli Enterprise Portal Server is the central repository for all user data.

Tivoli Enterprise Monitoring Server

The monitoring server acts as a collection and control point for alerts received from the monitoring agents, and collects their performance and availability data. The Tivoli Enterprise Monitoring Server is also a repository for historical data.

Tivoli Enterprise Monitoring Agent, Linux Kernel-based Virtual Machines agent

This monitoring agent collects data and distributes the data to the Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, Tivoli Enterprise Portal, Tivoli Data Warehouse, and IBM Dashboard Application Services Hub.

Multiple copies of this agent can run on the same system.

IBM Tivoli Netcool/OMNIBus

Tivoli Netcool/OMNIBus is an optional component and the recommended event management component. The Netcool/OMNIBus software is a service level management (SLM) system that delivers real-time, centralized monitoring of complex networks and IT domain events. Event information is tracked in a high-performance, in-memory database and presented to specific users through individually configurable filters and views. The software includes automation functions that you can use to perform intelligent processing on managed events. You can use this software to forward events for Tivoli Monitoring situations to Tivoli Netcool/OMNIBus.

IBM Tivoli Enterprise Console

The Tivoli Enterprise Console is an optional component that acts as a central collection point for events from various sources, including events from other Tivoli software applications, Tivoli partner applications, custom applications, network management platforms, and relational database systems. You can view these events through the Tivoli Enterprise Portal (by using the event viewer), and you can forward events from Tivoli Monitoring situations to the Tivoli Enterprise Console component. If you do not already use Tivoli Enterprise Console and need an event management component, you can choose to use IBM Tivoli Netcool/OMNIBus.

IBM Tivoli Common Reporting

Tivoli Common Reporting is a separately installable feature available to users of Tivoli software that provides a consistent approach to generating and customizing reports. Some individual products provide reports that are designed for use with Tivoli Common Reporting and have a consistent look and feel.

IBM Tivoli Application Dependency Discovery Manager (TADDM)

TADDM delivers automated discovery and configuration tracking capabilities to build application maps that provide real-time visibility into application complexity.

IBM Tivoli Business Service Manager

The Tivoli Business Service Manager component delivers real-time information to help you respond to alerts effectively based on business requirements. Optionally, you can use this component to meet service-level agreements (SLAs). Use the Tivoli Business Service Manager tools to help build a service model that you can integrate with Tivoli Netcool/OMNIBus alerts or optionally integrate with data from an SQL data source. Optional components provide access to data from other IBM Tivoli applications such as Tivoli Monitoring and TADDM.

IBM Dashboard Application Services Hub

The Dashboard Application Services Hub has a core set of components that provide such administrative essentials as network security and database management. This component replaces the Tivoli Integrated Portal component after version 2.2.

Agent Management Services

You can use IBM Tivoli Monitoring Agent Management Services to manage the Linux Kernel-based Virtual Machines agent.

Agent Management Services is available for the following IBM Tivoli Monitoring OS agents: Windows, Linux, and UNIX. The services are designed to keep the Linux Kernel-based Virtual Machines agent available, and to provide information about the status of the product to the Tivoli Enterprise Portal. IBM Tivoli Monitoring V6.2.2, Fix Pack 2 or later provides support for Agent Management Services. For more information about Agent Management Services, see *Agent Management Services* in the *IBM Tivoli Monitoring Administrator's Guide*.

User interface options

Installation of the base IBM Tivoli Monitoring software and other integrated applications provides various interfaces that you can use to work with your resources and data.

The following interfaces are available:

Tivoli Enterprise Portal user interface

You can run the Tivoli Enterprise Portal as a desktop application or a browser application. The client interface is a graphical user interface (GUI) based on Java on a Windows or Linux workstation. The browser application is automatically installed with the Tivoli Enterprise Portal Server. The desktop application is installed by using the Tivoli Monitoring installation media or with a Java Web Start application. To start the Tivoli Enterprise Portal browser client in your Internet browser, enter the URL for a specific Tivoli Enterprise Portal browser client installed on your Web server.

Command-line interface

You can use Tivoli Monitoring commands to manage the Tivoli Monitoring components and their configuration. You can also run commands at the Tivoli Enterprise Console event server or the Tivoli Netcool/OMNIBus ObjectServer to configure event synchronization for enterprise situations.

Manage Tivoli Enterprise Monitoring Services window

You can use the window for the Manage Tivoli Enterprise Monitoring Services utility to configure the agent and start Tivoli services not designated to start automatically.

IBM Tivoli Netcool/OMNIBus event list

You can use the Netcool/OMNIBus event list to monitor and manage events. An event is created when the Netcool/OMNIBus ObjectServer receives an event, alert, message, or data item. Each event is made up of columns (or fields) of information that are displayed in a row in the ObjectServer alerts.status table. The Tivoli Netcool/OMNIBus web GUI is also a web-based application that processes network events from one or more data sources and presents the event data in various graphical formats.

IBM Tivoli Enterprise Console

You can use the Tivoli Enterprise Console to help ensure the optimal availability of an IT service for an organization. The Tivoli Enterprise Console is an event management application that integrates system, network, database, and application management. If you do not already use Tivoli Enterprise Console and need an event management component, you can choose to use Tivoli Netcool/OMNIBus.

IBM Tivoli Common Reporting

Use the Tivoli Common Reporting web user interface for specifying report parameters and other report properties, generating formatted reports, scheduling reports, and viewing reports. This user interface is based on the Dashboard Application Services Hub.

IBM Tivoli Application Dependency Discovery Manager

The Discovery Management Console is the TADDM client user interface for managing discoveries.

IBM Tivoli Business Service Manager

The Tivoli Business Service Manager console provides a graphical user interface that you can use to logically link services and business requirements within the service model. The service model provides an operator with a second-by-second view of how an enterprise is performing at any moment in time or how the enterprise performed over a time period.

IBM Dashboard Application Services Hub

The Dashboard Application Services Hub provides an administrative console for applications that use this framework. It is a web-based console that provides common task navigation for products, aggregation of data from multiple products into a single view, and the passing of messages between views from different products. This interface replaces the Tivoli Integrated Portal component after version 2.2.

Chapter 2. Agent installation and configuration

Agent installation and configuration requires the use of the *IBM Tivoli Monitoring Installation and Setup Guide* and agent-specific installation and configuration information.

To install and configure the Linux Kernel-based Virtual Machines agent, use the *Installing monitoring agents* procedures in the *IBM Tivoli Monitoring Installation and Setup Guide* along with the agent-specific installation and configuration information.

If you are installing silently by using a response file, see *Performing a silent installation of IBM Tivoli Monitoring* in the *IBM Tivoli Monitoring Installation and Setup Guide*.

Requirements

Before installing and configuring the agent, make sure your environment meets the requirements for the IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines.

For information about system requirements, see the Software product compatibility reports (<http://publib.boulder.ibm.com/infocenter/prodguid/v1r0/clarify/index.html>). Search for the Tivoli Monitoring for Virtual Environments product.

Language pack installation

The steps for installing language packs depend on which operating system and mode of installation you are using.

To install a language pack for the agent support files on the Tivoli Enterprise Monitoring Server, the Tivoli Enterprise Monitoring Agent, and the Tivoli Enterprise Portal Server, make sure that you installed the product in the English language. Then use the steps for the operating system or mode of installation you are using:

- “Installing language packs on Windows systems”
- “Installing language packs on UNIX or Linux systems” on page 6
- “Silent installation of language packs on Windows, UNIX, or Linux systems” on page 6

Installing language packs on Windows systems

You can install the language packs on a Windows system.

Before you begin

First, make sure that you installed the product in the English language.

Procedure

1. On the language pack CD, double-click the `lpinstaller.bat` file to start the installation program.
2. Select the language of the installer and click **OK**.
3. In the Introduction panel, click **Next**
4. Click **Add/Update** and click **Next**.
5. Select the folder where the National Language Support package (NLSPackage) files are located. Typically, the NLSPackage files are located in the `nlspackage` folder where the installer executable file is located.

6. Select the language support for the agent of your choice and click **Next**. To make multiple selections, press Ctrl and select the language that you want.
7. Select the languages that you want to install and click **Next**.
8. Examine the installation summary page and click **Next** to begin installation.
9. After installation completes, click **Finish** to exit the installer.
10. Restart the Tivoli Enterprise Portal, Tivoli Enterprise Portal Server, and Eclipse Help Server if any of these components are installed.

Installing language packs on UNIX or Linux systems

You can install the language packs on a UNIX or Linux system.

Before you begin

First, make sure that you installed the product in the English language.

Procedure

1. Enter the `mkdir` command to create a temporary directory on the computer, for example, `mkdir dir_name`. Make sure that the full path of the directory does not contain any spaces.
2. Mount the language pack CD to the temporary directory that you created.
3. Enter the following command to start the installation program: `cd dir_name | installer.sh -c install_dir` where `install_dir` is where you installed IBM Tivoli Monitoring. Typically, the directory name is `/opt/IBM/ITM` for UNIX and Linux systems.
4. Select the language of the installer and click **OK**.
5. In the Introduction panel, click **Next**.
6. Click **Add/Update** and click **Next**.
7. Select the folder where the National Language Support package (NLSPackage) files are located. Typically, the NLSPackage files are located in the `nlspackage` folder where the installer executable file is located.
8. Select the language support for the agent of your choice and click **Next**. To make multiple selections, press Ctrl and select the language that you want.
9. Select the languages that you want to install and click **Next**.
10. Examine the installation summary page and click **Next** to begin installation.
11. After installation completes, click **Finish** to exit the installer.
12. Restart the Tivoli Enterprise Portal, Tivoli Enterprise Portal Server, and Eclipse Help Server if any of these components are installed.

Silent installation of language packs on Windows, UNIX, or Linux systems

You can use the silent-mode installation method to install the language packs. In silent mode, the installation process obtains the installation settings from a predefined response file. It does not prompt you for any information.

Before you begin

First, make sure that you installed the product in the English language.

Procedure

1. Copy and paste the `ITM_Agent_LP_silent.rsp` response file template as shown in "Response file example" on page 7.
2. Change the following parameter settings:

NLS_PACKAGE_FOLDER

Folder where the National Language Support package (NLSPackage) files are located. Typically, the NLSPackage files are located in the nlspackage folder, for example:
NLS_PACKAGE_FOLDER = //tmp//LP//nlspackage.

PROD_SELECTION_PKG

Name of the language pack to install. Several product components can be included in one language package. You might want to install only some of the available components in a language pack.

BASE_AGENT_FOUND_PKG_LIST

Agent for which you are installing language support. This value is usually the same as *PROD_SELECTION_PKG*.

LANG_SELECTION_LIST

Language you want to install.

3. Enter the command to install the language pack with a response file (silent installation):

- For Windows systems:
lpinstaller.bat -f *path_to_response_file*
- For UNIX or Linux systems:
lpinstaller.sh -c *candle_home* -f *path_to_response_file*

where *candle_home* is the IBM Tivoli Monitoring base directory.

Response file example

```
# IBM Tivoli Monitoring Agent Language Pack Silent Installation Operation
#
#This is a sample response file for silent installation mode for the IBM Tivoli
#Monitoring Common Language Pack Installer.
#.
#This file uses the IBM Tivoli Monitoring Common Agent Language Pack with the
#install package as an example.
#Note:
#This response file is for the INSTALLATION of language packs only.
#This file does not support UNINSTALLATION of language packs in silent mode.
#-----
#-----
#To successfully complete a silent installation of the the example of Common Agent
#localization pack, complete the following steps:
#
#1.Copy ITM_Agent_LP_silent.rsp to the directory where lpinstaller.bat or
#lpinstaller.sh is located (IBM Tivoli Monitoring Agent Language Pack build
#location).
#
#2.Modify the response file so that it is customized correctly and completely for
#your site.
# Complete all of the following steps in the response file.
#
#3.After customizing the response file, invoke the silent installation using the
#following command:
#For Windows:
# lpinstaller.bat -f <path_to_response_file>
#For UNIX and Linux:
# lpinstaller.sh -c <candle_home> -f <path_to_response_file>
#Note:<candle_home> is the IBM Tivoli Monitoring base directory.
#-----
#-----
#Force silent install mode.
#-----
INSTALLER_UI=silent
#-----
#Run add and update actions.
```

```

#-----
CHOSEN_INSTALL_SET=ADDUPD_SET
#-----
#NLS Package Folder, where the NLS Packages exist.
#For Windows:
# Use the backslash-backslash(\\) as a file separator (for example,
#C:\\zosgmv\\LCD7-3583-01\\nlspackage).
#For UNIX and Linux:
# Use the slash-slash (//) as a file separator (for example,
#//installtivo1i//lpsilenttest//nlspackage).
#-----
#NLS_PACKAGE_FOLDER=C:\\zosgmv\\LCD7-3583-01\\nlspackage
NLS_PACKAGE_FOLDER=//tmp//LP//nlspackage
#-----
#List the packages to process; both variables are required.
#Each variable requires that full paths are specified.
#Separate multiple entries with a semicolon (;).
#For Windows:
# Use the backslash-backslash(\\) as a file separator.
#For Unix and Linux:
# Use the slash-slash (//) as a file separator.
#-----
#PROD_SELECTION_PKG=C:\\zosgmv\\LCD7-3583-01\\nlspackage\\KIP-NLS.nlspkg
#BASE_AGENT_FOUND_PKG_LIST=C:\\zosgmv\\LCD7-3583-01\\nlspackage\\KIP-NLS.nlspkg
PROD_SELECTION_PKG=//tmp//LP//nlspackage//kex_nls.nlspkg;//tmp//LP//nlspackage//
koq_nls.nlspkg
BASE_AGENT_FOUND_PKG_LIST=//tmp//LP//nlspackage//kex_nls.nlspkg;//
tmp//LP//nlspackage//koq_nls.nlspkg
#-----
#List the languages to process.
#Separate multiple entries with semicolons.
#-----
LANG_SELECTION_LIST=pt_BR;fr;de;it;ja;ko;zh_CN;es;zh_TW

```

Agent-specific installation and configuration

In addition to the installation and configuration information in the *IBM Tivoli Monitoring Installation and Setup Guide*, use this agent-specific installation and configuration information to install the Linux Kernel-based Virtual Machines agent.

Virtualization hosts

You must take into account two aspects of your virtualization hosts when you configure the Linux Kernel-based Virtual Machines agent. The first aspect is how the hosts are organized: whether you have multiple kinds of virtualized loads and whether you want to migrate virtual machines between hosts. The second aspect is how you want to secure communications between the Linux Kernel-based Virtual Machines agent and the hypervisors.

The Linux Kernel-based Virtual Machines agent is both multi-instance and multiconnection:

- Multi-instance means that you can create multiple, independent copies of the Linux Kernel-based Virtual Machines agent.
- Multiconnection means that each instance of the Linux Kernel-based Virtual Machines agent can connect to and monitor multiple hypervisors. It is easier to compare the characteristics of several hypervisors when they are monitored by one agent instance.

One configuration strategy is to use a single agent instance to group together hosts with similar virtualized loads. By monitoring the supporting hosts with a single instance, you can more easily compare for each host the resource capacity and current resource consumption by virtualization. This type of monitoring can help you balance your virtualized workloads by migrating virtual machines from overburdened hosts to hosts that can support more work. Each kind of virtualized workload can be

monitored by its own Linux Kernel-based Virtual Machines agent. If you have instance-based monitoring, you can have one instance monitoring your web applications, another monitoring your virtual desktops, and so on.

The Linux Kernel-based Virtual Machines agent collects its metrics by connecting remotely to each libvirt hypervisor managing your QEMU-KVM virtual machines. The libvirt hypervisor can use several different remote transport protocols, as described on the Remote support page of the Libvirt Virtualization API website (<http://libvirt.org/remote.html>). The Linux Kernel-based Virtual Machines agent supports the SSH protocol, the TLS protocol, and the TCP protocol. While the SSH protocol and the TLS protocol provide production-level security, the use of the TCP protocol by the Linux Kernel-based Virtual Machines agent supports only an authentication of *none* and is intended for testing. Follow the instructions for implementing the SSH protocol, the TLS protocol, and the TCP protocol in the context of libvirt remote connections.

SSH protocol

For the SSH agent, assume you install the Linux Kernel-based Virtual Machines agent on Host A and you want to remotely monitor the hypervisor on Host B. First, you must configure the SSH agent so the SSH agent can make a connection from Host A to Host B without requiring you to include a password.

About this task

After configuration, you can start the Linux Kernel-based Virtual Machines agent and begin to monitor Host B. Assume you have several hosts that you want to monitor, for example, several Host Bs. For some helpful instructions on this topic, see the following procedure for accessing an SSH agent without a password at the Using the ssh-agent with ssh website (<http://mah.everybody.org/docs/ssh>).

Procedure

1. Log on to Host A with the same ID that will run the Linux Kernel-based Virtual Machines agent process, for example, the root user ID. Have available the ID on Host B that will be making the SSH connection, often also the root user ID.
2. Generate the `id_rsa` and `id_rsa.pub` keys on Host A. The keys are saved in `~/.ssh`:

```
$ ssh-keygen -t rsa
```
3. Copy the authorized keys from Host B, so you can add the public key for Host A to it:

```
$ scp Id on Host B@name or IP address of Host B:~/.ssh/authorized_keys ~/.ssh/authorized_keys_from_B
```
4. Append the public key for Host A to the end of the authorized keys for Host B:

```
cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys_from_B
```
5. Copy the authorized keys back to Host B. If you are monitoring multiple hosts, repeat steps 3, 4, and 5 for each host. You can remove `~/.ssh/authorized_keys_from_B` after this step:

```
$ scp ~/.ssh/authorized_keys_from_B Id on Host B@name or IP address of Host B:~/.ssh/authorized_keys
```
6. Add the following command to the `~/.bash_profile` of the current ID on Host A: **\$ eval `ssh-agent`**. Ensure you use the single back quotation mark (```), located under the tilde (`~`) on US keyboards, rather than the single quotation mark (`'`).
7. Add the identity to Host A. You are asked for the passphrase you used when the ID was created. If you receive the message "Could not open a connection to your authentication agent.", run the **exec ssh-agent bash** command (you can replace `bash` with the shell you are using) and then run the **ssh-add** command again:

```
$ ssh-add ~/.ssh/id_rsa
```
8. Test that the SSH agent can make a connection from Host A to Host B without entering the SSH password. If you are monitoring multiple hosts, test the connection for each host:

```
$ ssh Id on Host B@name or IP address of Host B
```

What to do next

When you have finished the configuration, check your work by using the **virsh** command by entering `virsh -c qemu+ssh://Id on Host B@name or IP address of Host B:port/system`

You can omit the **:port** section of the command if you have not changed the default SSH port. If the **virsh** command succeeds, the Linux Kernel-based Virtual Machines agent can connect.

Note: You must rerun the **ssh-add** command and supply the passphrase each time you restart Host A before you restart the Linux Kernel-based Virtual Machines agent on Host A. If you use SSH keychains, you can avoid having to reenter the passphrase. A discussion of SSH keychains is beyond the scope of this guide, but information is available on the Internet.

TLS protocol

TLS (Transport Layer Security) is often implemented in an organization to authenticate server-to-server communications. If you have already established TLS authentication on your servers, you can also use TLS for libvirt-to-libvirt communications.

About this task

For the TLS protocol, assume you install the Linux Kernel-based Virtual Machines agent on Host A and you want to remotely monitor the hypervisor on Host B.

Procedure

1. Log in to Host B and confirm you have installed the `gnutls` and `gnutls-utils` packages.
2. Edit `/etc/libvirt/libvirtd.conf` to make sure that `listen_tls` is enabled and the `tls_port` is 16514 (the default).
3. Go to libvirt.org/remote.html and follow the instructions for setting up a certificate authority between Host A and Host B. Pay special attention to the sections of Setting up a Certificate Authority (CA), Issuing server certificates, and Issuing client certificates.
4. Restart the libvirt daemon on Host B in listening mode by running it with the `--listen` flag or by editing `/etc/sysconfig/libvirtd` and uncommenting the `LIBVIRT_ARGS="--listen"` line.

What to do next

When you have finished the configuration, check your work using the **virsh** command by entering `virsh -c qemu+tls://name or IP address of Host B:port/system`. You can omit the **:port** section of the command if you have not changed the default TLS port. If the **virsh** command succeeds, the Linux Kernel-based Virtual Machines agent can connect.

TCP protocol

Use the TCP protocol only for testing.

About this task

For TCP, assume you install the Linux Kernel-based Virtual Machines agent on Host A and you want to remotely monitor the hypervisor on Host B. Follow these steps:

Procedure

1. Log in to Host B.
2. Edit `/etc/libvirt/libvirtd.conf` to make sure that `listen_tcp` is enabled and `tcp_port` is 16509 (the default).
3. Edit `/etc/libvirt/libvirtd.conf` to set `auth_tcp` to "none". This step instructs TCP not to perform any authentication.

- Restart the libvirt daemon on Host B in listening mode by running it with the `--listen` flag or by editing `/etc/sysconfig/libvirtd` and uncommenting the `LIBVIRT_ARGS="--listen"` line.

What to do next

When you have finished the configuration, check your work using the `virsh` command by entering `virsh -c qemu+tcp://name or IP address of Host B:port/system`. You can omit the `:port` section of the command if you have not changed the default TCP port. If the `virsh` command succeeds, the Linux Kernel-based Virtual Machines agent can connect.

Configuration values

For both local and remote configuration, you provide the configuration values for the agent to operate.

When you are configuring an agent, a panel is displayed so you can enter each value. When a default value exists, this value is pre-entered into the field. If a field represents a password, two entry fields are displayed. You must enter the same value in each field. The values you type are not displayed to help maintain the security of these values.

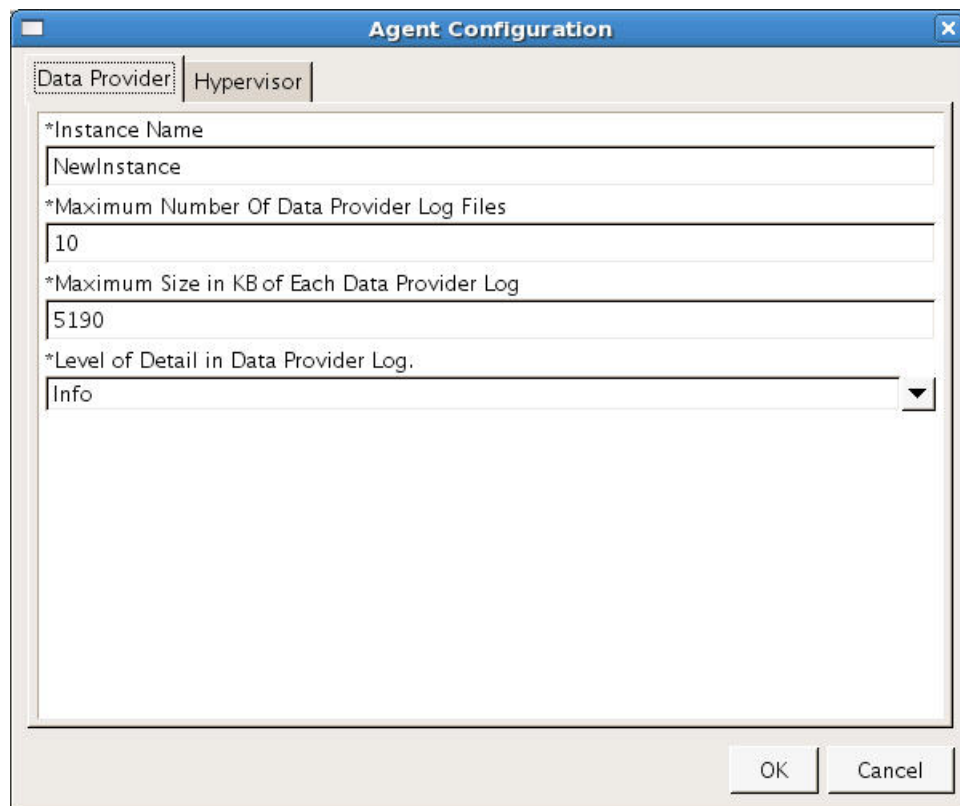
The configuration for this agent is organized into the following groups:

Data Provider (DATA_PROVIDER)

This section provides the logging characteristics that the data provider uses.

The configuration elements defined in this group are always present in the configuration of the agent.

This group defines information that applies to the entire agent.



The screenshot shows a dialog box titled "Agent Configuration" with two tabs: "Data Provider" (selected) and "Hypervisor". The "Data Provider" tab contains the following fields:

- *Instance Name: NewInstance
- *Maximum Number Of Data Provider Log Files: 10
- *Maximum Size in KB of Each Data Provider Log: 5190
- *Level of Detail in Data Provider Log: Info (dropdown menu)

At the bottom right of the dialog box are "OK" and "Cancel" buttons.

Figure 1. Agent Configuration: Data Provider

Maximum Number Of Data Provider Log Files (KV1_LOG_FILE_MAX_COUNT)

This is the number of log files that is created before rolling over.

The type is numeric.

This value is required.

Default value: 10

Maximum Size in KB of Each Data Provider Log (KV1_LOG_FILE_MAX_SIZE)

This value is the maximum size in KB that a log file reaches before moving to the next log file.

The type is numeric.

This value is required.

Default value: 5190

Level of Detail in Data Provider Log. (KV1_LOG_LEVEL)

This value controls how many log messages the agent writes and at what level of detail.

The type is one of the following values: **Off, Severe, Warning, Info, Fine, Finer, Finest, All.**

This value is required.

Default value: INFO

Hypervisor (HYPERVISOR)

This section provides the connection information for each hypervisor being monitored.

The configuration elements defined in this group are always present in the agent's configuration.

Use the information in this group to create additional subnodes.

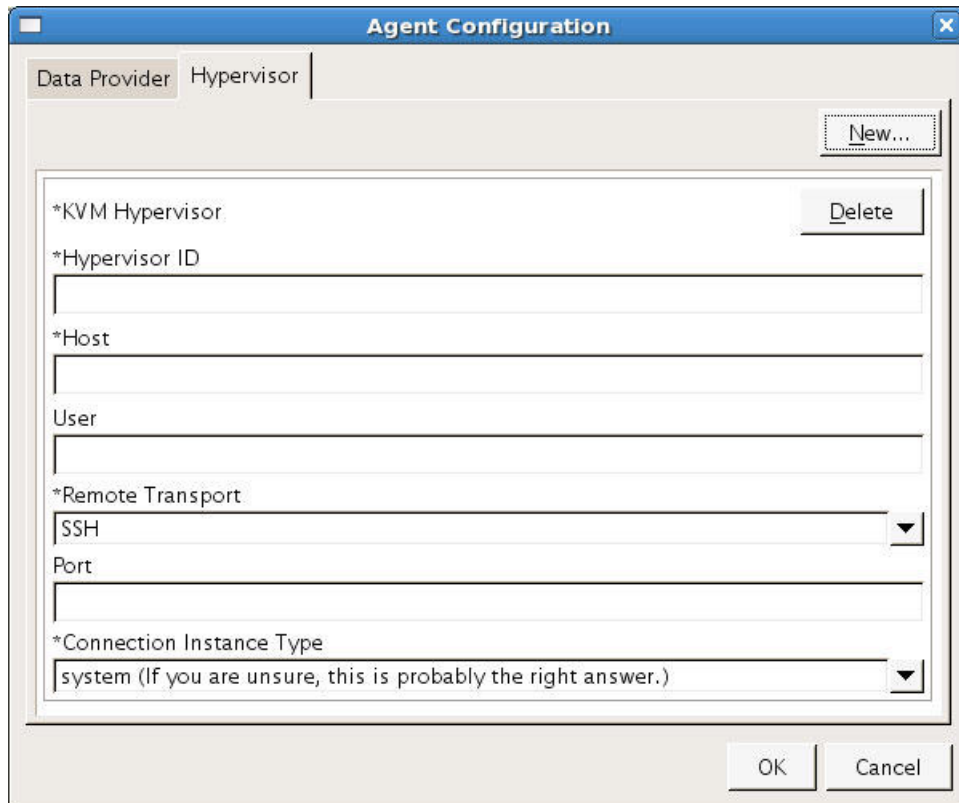


Figure 2. Agent Configuration: Hypervisor

Hypervisor ID (Hypervisor ID)

A unique identifier for this hypervisor.

The type is string.

This value is required.

Default value: None

Host (HOST_ADDRESS)

The host name or IP address of the KVM hypervisor.

The type is string.

This value is required.

Default value: None

User (USERNAME)

A user name on the KVM hypervisor that has sufficient privileges to collect monitoring data. It is only needed for SSH transport.

The type is string.

This value is optional.

Default value: None

Remote Transport (PROTOCOL)

This value controls which protocol the local libvirt uses to connect to remote libvirts.

The type is one of the following values: **SSH, TLS, TCP (Unencrypted -- not recommended for production use).**

This value is required.

Default value: ssh

Port (PORT)

The port used by the transport protocol to make the libvirt connection. It is only needed if the standard ports have been changed (22 for SSH, 16514 for TLS, 16509 for TCP).

The type is numeric.

This value is optional.

Default value: None

Connection Instance Type (CONNECTION_MODE)

This value controls whether the local libvirt connects to the privileged system driver or the per-user unprivilege session driver.

The type is one of the following values:

- system

Note: If you are unsure, this answer is probably correct.

- session

This value is required.

Default value: system

Remote installation and configuration

You can install the monitoring agent remotely from the Tivoli Enterprise Portal or from the command line.

When installing the agent remotely, you must provide the configuration values for the agent to operate. See “Configuration values” on page 11.

To install from the portal, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

To remotely install or configure an agent through the Tivoli Enterprise Portal, you must have installed the application support for that agent (Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, and Tivoli Enterprise Portal). You must also have installed the agent bundle into the Remote Deploy Depot.

For information about displaying the configuration options that are available to use with the **configureSystem** or **addSystem** commands see "tacmd describeSystemType" in the *IBM Tivoli Monitoring Command Reference*.

If you are using the command line, the following command is an example of remote installation and configuration for Windows operating systems:

```
tacmd addsystem -t v1 -n Primary:sample.node.name:LZ
-p DATA_PROVIDER.KV1_LOG_FILE_MAX_COUNT=10
DATA_PROVIDER.KV1_LOG_FILE_MAX_SIZE=5190
DATA_PROVIDER.KV1_LOG_LEVEL=INFO
INSTANCE=instanceName
HYPERVISOR:UniqueDataSourceID.HOST_ADDRESS=hostAddressOrName
HYPERVISOR:UniqueDataSourceID.USERNAME=user
HYPERVISOR:UniqueDataSourceID.PROTOCOL=(ssh|tls|tcp)
HYPERVISOR:UniqueDataSourceID.PORT=port
HYPERVISOR:UniqueDataSourceID.CONNECTION_MODE=system
```

Chapter 3. Workspaces reference

A workspace is the working area of the Tivoli Enterprise Portal application window. The Navigator tree contains a list of the workspaces provided by the agent.

About workspaces

Use the Navigator tree to select the workspace you want to see. As part of the application window, the status bar shows the Tivoli Enterprise Portal Server name and port number to which the displayed information applies and the ID of the current user.

When you select an item in the Navigator tree, a default workspace is displayed. When you right-click a Navigator item, a menu that includes a Workspace item is displayed. The Workspace item contains a list of workspaces for that Navigator item. Each workspace has at least one view. Some views have links to other workspaces. You can also use the Workspace Gallery tool as described in the *Tivoli Enterprise Portal User's Guide* to open workspaces.

The workspaces in the Navigator are displayed in a Physical view that shows your enterprise as a physical mapping or a dynamically populated logical view that is agent-specific. You can also create a Logical view. The Physical view is the default view.

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.

Workspace views can be any combination of query-based views, event views, and special purpose views.

Additional information about workspaces

For more information about creating, customizing, and working with workspaces, see "Using workspaces" in the *Tivoli Enterprise Portal User's Guide*.

For a list of the predefined workspaces for this monitoring agent and a description of each workspace, see Predefined workspaces and the information about each individual workspace.

Some attribute groups for this monitoring agent might not be represented in the predefined workspaces or views for this agent. For a full list of the attribute groups, see "Attribute groups for the monitoring agent" on page 21.

Predefined workspaces

The Linux Kernel-based Virtual Machines agent provides predefined workspaces, which are organized by Navigator item.

- Linux Kernel-based Virtual Machines Navigator item
 - Linux Kernel-based Virtual Machines workspace
 - Host CPU Comparison workspace
 - Host Memory Allocation Comparison workspace
 - Host Memory Use Comparison workspace
- Host Navigator item
 - Host workspace
 - Host Detail workspace

- Host Overview workspace
- Host Relations workspace
- Storage Pool Navigator item
 - Storage Pool workspace
 - Storage Pool Detail workspace
 - Storage Pool Overview workspace
- Virtual Machine Navigator item
 - Virtual Machine workspace
 - Virtual Machine Detail workspace
 - Virtual Machine Overview workspace

Workspace descriptions

Each workspace description provides information about the workspace such as the purpose and a list of views in the workspace.

Workspaces are listed under Navigator items.

Linux Kernel-based Virtual Machines Navigator item

The workspace descriptions are organized by the Navigator item to which the workspaces are relevant.

Linux Kernel-based Virtual Machines workspace

These workspaces show the most heavily used hosts and the least heavily used hosts based on various resources. The hosts are compared in absolute measurements and in percentage measurements.

This workspace contains the following views:

Host CPU Comparison

This workspace shows the most heavily used hosts and the least heavily used hosts in terms of CPU use.

Host Memory Allocation Comparison

This workspace shows the most heavily used hosts and the least heavily used hosts in terms of memory allocation.

Host Memory Use Comparison

This workspace shows the most heavily used hosts and the least heavily used hosts in terms of memory use.

Host CPU Comparison workspace

This workspace shows the most heavily used hosts and the least heavily used hosts in terms of CPU use. The hosts are compared in absolute measurements and in percentage measurements. The absolute measurements use CPU GHz to enable comparison between different machine types. All the CPU measurements can be seen by expanding the table view that holds all the detailed data.

This workspace contains the following views:

Top VM CPU GHz By Host

This view shows the most heavily used hosts based on the absolute measurement of CPU GHz consumed by the virtual machines on each host.

Top Unaccounted CPU GHz By Host

This view shows the least heavily used hosts based on the absolute measurement of CPU GHz unaccounted for by the virtual machines consumption on each host.

Top VM Consumed CPU Percent By Host

This view shows the most heavily used hosts based on the percentage of the CPU consumed by the virtual machines from the number of CPUs on each host.

Top Unaccounted CPU Percent By Host

This view shows the least heavily used hosts based on the percentage of the CPU that remains unaccounted for after subtracting the CPU consumed by the virtual machines on each host.

Hosts This view shows all the CPU measurements for each host. The link at the beginning of each row takes you to the Host Detail workspace for the selected host.

Host Memory Allocation Comparison workspace

This workspace shows the most heavily used hosts and the least heavily used hosts in terms of memory allocation. The hosts are compared in absolute measurements and in percentage measurements. The absolute measurements use GB of memory. All the memory measurements can be seen by expanding the table view that holds all the detailed data.

This workspace contains the following views:

Top Allocated Memory By Host

This view shows the most heavily used hosts based on the absolute measurement of GB of memory allocated to the virtual machines on each host.

Top Unallocated Memory By Host

This view shows the least heavily used hosts based on the absolute measurement of GB of memory allocated to the virtual machines on each host.

Top Allocated Memory Percent By Host

This view shows the most heavily used hosts based on the percentage of the GB of memory allocated to the virtual machines from the GB of memory on each host.

Top Unallocated Memory Percent By Host

This view shows the least heavily used hosts based on the percentage of the memory that remains unallocated after subtracting the memory allocated to the virtual machines on each host.

Hosts This view shows all the memory measurements for each host. The link at the beginning of each row takes you to the Host Detail workspace for the selected host.

Host Memory Use Comparison workspace

This workspace shows the most heavily used hosts and the least heavily used hosts in terms of memory use. The hosts are compared in absolute measurements and in percentage measurements. The absolute measurements use GB of memory. All the memory measurements can be seen by expanding the table view that holds all the detailed data.

This workspace contains the following views:

Top Used Memory By Host

This view shows the most heavily used hosts based on the absolute measurement of GB of memory used by the virtual machines on each host.

Top Unused Memory By Host

This view shows the least heavily used hosts based on the absolute measurement of GB of memory used by the virtual machines on each host.

Top Used Memory Percent By Host

This view shows the most heavily used hosts based on the percentage of the GB of memory used by the virtual machines from the GB of memory on each host.

Top Unused Memory Percent By Host

This view shows the least heavily used hosts based on the percentage of the memory that remains unused after subtracting the memory used by the virtual machines on each host.

Hosts This view shows all the memory measurements for each host. The link at the beginning of each row takes you to the Host Detail workspace for the selected host.

Host Navigator item

The workspace descriptions are organized by the Navigator item to which the workspaces are relevant.

Host workspace

These workspaces show all the hosts that are being monitored by an agent instance, and detailed information about individual hosts and the data center resources related to a host.

This workspace contains the following views:

Host Overview

This workspace shows all the hosts that are being monitored by an agent instance.

Host Detail

This workspace shows detailed information about one host, including trend graphs of CPU and memory use, as well as defining characteristics, and CPU and memory metrics.

Host Relations

This workspace shows the data center resources that are related to one host.

Host Detail workspace

This view shows a trend graph of CPU capacity allocated and consumed by the virtual machines on the host.

This workspace contains the following views:

CPU Allocated and Consumed

This view shows a trend graph of CPU capacity allocated and consumed by the virtual machines on the host.

Memory Allocated and Consumed

This view shows a trend graph of memory allocated-to and used-by the virtual machines on the host.

Host Definition

This view shows the defining characteristics of the host.

Host CPU

This view shows the CPU metrics of the host.

Host Memory

This view shows the memory metrics of the host.

Host Overview workspace

This workspace shows all the hosts that are being monitored by an agent instance. Each view shows a different set of metrics about every host. Every row has a link that takes you to the Host Detail workspace for the selected host.

This workspace contains the following views:

Host CPU Metrics

This view shows the CPU metrics of every host that is being monitored by an agent instance.

Host Memory Metrics

This view shows the memory metrics of every host that is being monitored by an agent instance.

Host Definitions

This view shows the defining characteristics of every host that is being monitored by an agent instance.

Host Relations workspace

This workspace shows the data center resources that are related to one host.

This workspace contains the following views:

Virtual Machines

This view shows the virtual machines on the host.

Storage Pools

This view shows the storage pools on the host.

Storage Pool Navigator item

The workspace descriptions are organized by the Navigator item to which the workspaces are relevant.

Storage Pool workspace

These workspaces show an overview of the most heavily used and least heavily used storage pools monitored by an agent instance, as well as detailed information about individual storage pools.

This workspace contains the following views:

Storage Pool Overview

This workspace shows an overview of the most heavily used and least heavily used storage pools monitored by an agent instance.

Storage Pool Detail

This workspace shows detailed information about one storage pool, including trend graphs of percent storage used and GB of storage available, as well as defining characteristics of the storage pool.

Storage Pool Detail workspace

This workspace shows detailed information about one storage pool, including trend graphs of percent storage used and GB of storage available, as well as defining characteristics of the storage pool.

This workspace contains the following views:

Percent Used

This view shows a trend graph of percent used of the storage pool.

Storage Available (GB)

This view shows a trend graph of the GB of storage available on the storage pool.

Storage Pool Detail

This view shows the detail metrics of the storage pool.

Storage Pool Overview workspace

This workspace shows an overview of the most heavily used and least heavily used storage pools monitored by an agent instance.

This workspace contains the following views:

Percent Used

This view shows the most heavily used storage pools based on percent used.

Storage Available (GB)

This view shows the least heavily used storage pools based on GB of storage available.

Storage Pools

This view shows the detail metrics of every storage pool monitored by an agent instance.

Virtual Machine Navigator item

The workspace descriptions are organized by the Navigator item to which the workspaces are relevant.

Virtual Machine workspace

These workspaces show an overview of the virtual machines monitored by an agent instance, including the virtual machines consuming the most CPU and the most memory, as well as detailed information about individual virtual machines.

This workspace contains the following views:

Virtual Machine Overview

This workspace shows an overview of the virtual machines monitored by an agent instance, including the virtual machines consuming the most CPU and the most memory.

Virtual Machine Detail

This workspace shows detailed information about one virtual machine, including trend graphs of CPU and memory use.

Virtual Machine Detail workspace

This workspace shows detailed information about one virtual machine, including trend graphs of CPU and memory use.

This workspace contains the following views:

CPUs Consumed

This view shows a trend graph of the number of CPUs consumed by the virtual machine.

CPU Percent

This view shows a trend graph of the number of CPUs consumed as a percent of the number of CPUs allocated to the virtual machine.

Memory Used

This view shows a trend graph of the GB of memory used by the virtual machine.

Memory Percent

This view shows a trend graph of amount of memory used as a percent of the memory allocated to the virtual machine.

Virtual Machine Detail

This view shows the detail metrics of the virtual machine.

Virtual Machine Overview workspace

This workspace shows an overview of the virtual machines monitored by an agent instance, including the virtual machines consuming the most CPU and the most memory.

This workspace contains the following views:

Top VM CPUs Consumed

This view shows the virtual machines that are consuming the most CPU.

Top VM Memory Used (GB)

This view shows the virtual machines that are consuming the most memory.

Virtual Machines

This view shows the detail metrics of every virtual machine monitored by an agent instance.

Chapter 4. Attributes reference

Attributes are the application properties that are being measured and reported by the IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines.

About attributes

Attributes are organized into attribute groups. Attributes in an attribute group relate to a single object such as an application, or to a single kind of data such as status information.

Attributes in a group can be used in queries, query-based views, situations, policy workflows, take action definitions, and launch application definitions. Chart or table views and situations are two examples of how attributes in a group can be used:

- 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 Properties editor to 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 values you assign to the situation attributes are compared with the values collected by the Linux Kernel-based Virtual Machines agent and registers an *event* if the condition is met. You are alerted to events by indicator icons that are displayed in the Navigator.

Additional information about attributes

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

For a list of the attribute groups, a list of the attributes in each attribute group, and descriptions of the attributes for this monitoring agent, see "Attribute groups for the monitoring agent" and "Attributes in each attribute group" on page 22.

Attribute groups for the monitoring agent

The Linux Kernel-based Virtual Machines agent contains the following attribute groups. The table name depends on the maximum table name limits of the target database being used for the Tivoli Data Warehouse. If the maximum name is 30 characters, any warehouse table name longer than 30 characters is shortened to 30 characters.

- Attribute group name: Host CPU
 - Table name: KV1HOSTCG
 - Warehouse table name: KV1_HOST_CPU or KV1HOSTCG
- Attribute group name: Host Memory
 - Table name: KV1HOSTMG
 - Warehouse table name: KV1_HOST_MEMORY or KV1HOSTMG
- Attribute group name: Hosts
 - Table name: KV1HOSTAG
 - Warehouse table name: KV1_HOSTS or KV1HOSTAG
- Attribute group name: Performance Object Status

- Table name: KV1POBJST
- Warehouse table name: KV1_PERFORMANCE_OBJECT_STATUS or KV1POBJST
- Attribute group name: Scheduler Parameters
 - Table name: KV1SCHPAG
 - Warehouse table name: KV1_SCHEDULER_PARAMETERS or KV1SCHPAG
- Attribute group name: Storage Pools
 - Table name: KV1STGPLAG
 - Warehouse table name: KV1_STORAGE_POOLS or KV1STGPLAG
- Attribute group name: Virtual Machines
 - Table name: KV1VMACHAG
 - Warehouse table name: KV1_VIRTUAL_MACHINES or KV1VMACHAG

Attributes in each attribute group

Attributes in each Linux Kernel-based Virtual Machines agent attribute group collect data that the agent uses for monitoring.

The descriptions of the attribute groups contain the following information:

Historical group

Whether the attribute group is a historical type that you can roll off to a data warehouse.

Attribute descriptions

Information such as description, type, source, and warehouse name, as applicable, for each attribute in the attribute group.

Some attributes are designated as key attributes. A *key attribute* is an attribute that is used in warehouse aggregation to identify rows of data that represent the same object.

The Source information sometimes uses C programming code syntax for if-then-else clauses to describe how an attribute is derived, for example:

```
(CPU_Pct < 0 ) || (Memory_Pct < 0 )? 0 : 1
```

This example means that if the CPU_Pct attribute is less than 0 or if the Memory_Pct attribute is less than 0, then the attribute is set to 0. Otherwise, the attribute is set to 1.

Host CPU attribute group

The Host CPU attribute group contains information about host CPU capacity and virtual machine CPU consumption.

Historical group

This attribute group is eligible for use with Tivoli Data Warehouse.

Attribute descriptions

The following list contains information about each attribute in the Host CPU attribute group:

Node attribute: This attribute is a key attribute.

Description

The managed system name of the agent.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

NODE

Timestamp attribute

Description

The local time at the agent when the data was collected.

Type
String

Source
The source for this attribute is the agent.

Warehouse name
TIMESTAMP

Host Name attribute: This attribute is a key attribute.

Description
The name of the host.

Type
String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name
HOST_NAME

Number of CPUs attribute

Description
The number of CPUs in the host.

Type
Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name
NUMBER_OF_CPUS or CPU_COUNT

Number of CPU GHz attribute

Description
The number of CPU GHz in the host.

Type
Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name
NUMBER_OF_CPU_GHZ or GHZ_COUNT

VM CPUs Consumed attribute

Description
The number of CPUs consumed by the Virtual Machines.

Type
Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VM_CPUS_CONSUMED or CPU_CNSMD

VM CPU GHz Consumed attribute

Description

The number of CPU GHz consumed by the Virtual Machines.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VM_CPU_GHZ_CONSUMED or GHZ_CNSMD

CPUs Unaccounted For attribute

Description

The number of CPU unaccounted for; includes system overhead and idle capacity.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPUS_UNACCOUNTED_FOR or CPU_UNACCT

CPU GHz Unaccounted For attribute

Description

The number of CPU GHz unaccounted for; includes system overhead and idle capacity.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_GHZ_UNACCOUNTED_FOR or GHZ_UNACCT

CPUs Consumed Per VM attribute

Description

The number of CPUs consumed by the average virtual machine.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_CONSUMED_PER_VM or CPU_PER_VM

CPU GHz Consumed Per VM attribute

Description

The number of CPU GHz consumed by the average virtual machine.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_GHZ_CONSUMED_PER_VM or GHZ_PER_VM

Number of Active VMs attribute

Description

The number of active virtual machines.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

NUMBER_OF_ACTIVE_VMS or ACTV_COUNT

CPU Percent Consumed by VMs attribute

Description

The very approximate percent consumed based on CPU capacity for the host and the consumption by the VMs. Does not take into account hypervisor overhead or non-virtual machine processes.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_PERCENT_CONSUMED or CNSMD_P

CPU Percent Unaccounted attribute

Description

The very approximate percent available based on CPU capacity for the host. Does not take into account hypervisor overhead or non-virtual machine processes.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_PERCENT_UNACCOUNTED or UNACCT_P

VM CPU Fit Estimate attribute**Description**

The very approximate number of virtual machines that fit in the unaccounted CPU capacity.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VM_CPU_FIT_ESTIMATE or CPU_FIT

Virtual CPUs Allocated attribute**Description**

The number of CPUs allocated to active Virtual Machines.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VIRTUAL_CPUS_ALLOCATED or CPU_ALLO

Over Under Allocated CPU attribute**Description**

The number of CPUs allocated to active virtual machines over or under the amount available of CPUs. Negative values indicate over allocation.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

OVER_UNDER_ALLOCATED_CPU or OV_UN_CPU

Over Under Allocated CPU GHz attribute**Description**

The amount of CPU GHz allocated to active virtual machines over or under the amount available of CPU GHz. Negative values indicate over allocation.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

OVER_UNDER_ALLOCATED_CPU_GHZ or OV_UN_GHZ

Percent CPU Allocated to VMs attribute**Description**

The percent of CPU that is allocated to virtual machines.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

PERCENT_CPU_ALLOCATED or C_ALLO_P

Host Memory attribute group

The Host Memory attribute group contains information about host memory capacity and virtual machine memory consumption.

Historical group

This attribute group is eligible for use with Tivoli Data Warehouse.

Attribute descriptions

The following list contains information about each attribute in the Host Memory attribute group:

Node attribute: This attribute is a key attribute.

Description

The managed system name of the agent.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

NODE

Timestamp attribute**Description**

The local time at the agent when the data was collected.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

TIMESTAMP

Host Name attribute: This attribute is a key attribute.

Description

The name of the host.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

HOST_NAME

Memory Size attribute

Description

The size of main memory in the host in GB

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

MEMORY_SIZE or MEMRY_SIZE

VM Memory Allocated attribute**Description**

The sum of the amounts of memory promised to each virtual machine in GB

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VM_MEMORY_ALLOCATED or MEMRY_ALLO

VM Memory Used attribute**Description**

The sum of the amounts of memory used by each virtual machine in GB

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VM_MEMORY_USED or MEMRY_USED

Over Under Allocated attribute**Description**

The amount of memory allocated to virtual machines over or under the amount available in GB; negative values indicate over allocation.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

OVER_UNDER_ALLOCATED or OV_UN_ALLO

Over Under Used attribute

Description

The amount of memory used by the virtual machines over or under the amount available in GB; negative values indicate over use.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

OVER_UNDER_USED or OV_UN_USED

Percent Memory Allocated attribute**Description**

The percent of memory that is allocated to virtual machines

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

PERCENT_MEMORY_ALLOCATED or M_ALLO_P

Percent Memory Used attribute**Description**

The percent of memory that is used by virtual machines

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

PERCENT_MEMORY_USED or M_USED_P

Percent Memory Unallocated attribute**Description**

The percent of memory that is unallocated to virtual machines

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

PERCENT_MEMORY_UNALLOCATED or M_UNALLO_P

Percent Memory Unused attribute

Description

The percent of memory that is unused by virtual machines

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

PERCENT_MEMORY_UNUSED or M_UNUSED_P

Number of Active VMs attribute**Description**

The number of active virtual machines

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

NUMBER_OF_ACTIVE_VMS or ACTV_COUNT

Memory Allocated Per VM attribute**Description**

The memory allocated to an average virtual machine

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

MEMORY_ALLOCATED_PER_VM or ALC_PER_VM

Memory Used Per VM attribute**Description**

The memory used by the average virtual machine

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

MEMORY_USED_PER_VM or USD_PER_VM

VM Memory Allocated Fit Estimate attribute**Description**

The very approximate number of virtual machines that fit in the unallocated memory, based on the average memory allocated per virtual machine.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VM_MEMORY_ALLOCATED_FIT_ESTIMATE or M_ALLO_FIT

VM Memory Used Fit Estimate attribute**Description**

The very approximate number of virtual machines that fit in the unused memory, based on the average memory used per virtual machine.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VM_MEMORY_USED_FIT_ESTIMATE or M_USED_FIT

Hosts attribute group

The host attribute group contains information about the host: name and address, hypervisor URI and protocol, number and configuration of CPUs, size of memory, and number of virtual machines.

Historical group

This attribute group is eligible for use with Tivoli Data Warehouse.

Attribute descriptions

The following list contains information about each attribute in the Hosts attribute group:

Node attribute: This attribute is a key attribute.

Description

The managed system name of the agent.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

NODE

Timestamp attribute**Description**

The local time at the agent when the data was collected.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

TIMESTAMP

Host Name attribute: This attribute is a key attribute.

Description

The name of the host.

Type
String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name
HOST_NAME

IP Address attribute

Description
The IP address used to create the Hypervisor URI.

Type
String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name
IP_ADDRESS or IP_ADDR

Hypervisor URI attribute

Description
The URI used for the connection to the hypervisor.

Type
String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name
HYPERVISOR_URI or HYPER_URI

Protocol attribute

Description
The communications protocol used by the connection to the hypervisor.

Type
String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name
PROTOCOL

CPU Model attribute

Description
The model of the CPUs in the host.

Type
String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_MODEL

Number of CPUs attribute

Description

The number of CPUs in the host.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

NUMBER_OF_CPUS or CPU_COUNT

CPU Frequency (GHz) attribute

Description

The frequency of the CPUs in the host in GHz.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_FREQUENCY or CPU_FREQ

Number of CPU GHz attribute

Description

The number of CPU GHz in the host.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

NUMBER_OF_CPU_GHZ or GHZ_COUNT

Number of Nodes attribute

Description

The number of NUMA cells in the host; 1 means uniform memory access.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

NUMBER_OF_NODES or NODE_COUNT

Sockets Per Node attribute

Description

The number of CPU sockets per node in the host.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

SOCKETS_PER_NODE or SKTS_PER

Cores Per Socket attribute

Description

The number of cores per socket in this host.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CORES_PER_SOCKET or CORES_PER

Threads Per Core attribute

Description

The number of threads per core in the host.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

THREADS_PER_CORE or THRDS_PER

Memory Size (GB) attribute

Description

The size of main memory in the host in GB.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

MEMORY_SIZE or MEMRY_SIZE

Number of VMs attribute

Description

The number of defined virtual machines in the host.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the

Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

NUMBER_OF_VMS or VM_COUNT

Number of Active VMs attribute

Description

The number of active virtual machines in the host.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

NUMBER_OF_ACTIVE_VMS or ACTV_COUNT

Performance Object Status attribute group

The Performance Object Status attribute group contains information that reflects the status of other attribute groups so you can see the status of all of the performance objects that make up this application all at once. Each of these other performance attribute groups is represented by a row in this table (or other type of view). The status for an attribute group reflects the result of the last attempt to collect data for that attribute group, which allows you to see whether the agent is performing correctly. Unlike other attribute groups, the Performance Object Status attribute group does not reflect the state of the monitored application. This attribute group is most often used to determine why data is not available for one of the performance attribute groups.

Historical group

This attribute group is eligible for use with Tivoli Data Warehouse.

Attribute descriptions

The following list contains information about each attribute in the Performance Object Status attribute group:

Node attribute: This attribute is a key attribute.

Description

The managed system name of the agent.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

NODE

Timestamp attribute

Description

The local time at the agent when the data was collected.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

TIMESTAMP

Query Name attribute: This attribute is a key attribute.

Description

The name of the attribute group.

Type

String

Warehouse name

QUERY_NAME or ATTRGRP

Object Name attribute**Description**

The name of the performance object.

Type

String

Warehouse name

OBJECT_NAME or OBJNAME

Object Type attribute**Description**

The type of the performance object.

Type

Integer with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- WMI (0)
- PERFMON (1)
- WMI ASSOCIATION GROUP (2)
- JMX (3)
- SNMP (4)
- SHELL COMMAND (5)
- JOINED GROUPS (6)
- CIMOM (7)
- CUSTOM (8)
- ROLLUP DATA (9)
- WMI REMOTE DATA (10)
- LOG FILE (11)
- JDBC (12)
- CONFIG DISCOVERY (13)
- NT EVENT LOG (14)
- FILTER (15)
- SNMP EVENT (16)
- PING (17)
- DIRECTOR DATA (18)
- DIRECTOR EVENT (19)
- SSH REMOTE SHELL COMMAND (20)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

OBJECT_TYPE or OBJTYPE

Object Status attribute**Description**

The status of the performance object.

Type

Integer with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- ACTIVE (0)
- INACTIVE (1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

OBJECT_STATUS or OBJSTTS

Error Code attribute

Description

The error code that is associated with the query.

Type

Integer with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- NO ERROR (0)
- GENERAL ERROR (1)
- OBJECT NOT FOUND (2)
- COUNTER NOT FOUND (3)
- NAMESPACE ERROR (4)
- OBJECT CURRENTLY UNAVAILABLE (5)
- COM LIBRARY INIT FAILURE (6)
- SECURITY INIT FAILURE (7)
- PROXY SECURITY FAILURE (9)
- NO INSTANCES RETURNED (10)
- ASSOCIATOR QUERY FAILED (11)
- REFERENCE QUERY FAILED (12)
- NO RESPONSE RECEIVED (13)
- CANNOT FIND JOINED QUERY (14)
- CANNOT FIND JOIN ATTRIBUTE IN QUERY 1 RESULTS (15)
- CANNOT FIND JOIN ATTRIBUTE IN QUERY 2 RESULTS (16)
- QUERY 1 NOT A SINGLETON (17)
- QUERY 2 NOT A SINGLETON (18)
- NO INSTANCES RETURNED IN QUERY 1 (19)
- NO INSTANCES RETURNED IN QUERY 2 (20)
- CANNOT FIND ROLLUP QUERY (21)
- CANNOT FIND ROLLUP ATTRIBUTE (22)
- FILE OFFLINE (23)
- NO HOSTNAME (24)
- MISSING LIBRARY (25)
- ATTRIBUTE COUNT MISMATCH (26)
- ATTRIBUTE NAME MISMATCH (27)
- COMMON DATA PROVIDER NOT STARTED (28)
- CALLBACK REGISTRATION ERROR (29)
- MDL LOAD ERROR (30)
- AUTHENTICATION FAILED (31)
- CANNOT RESOLVE HOST NAME (32)
- SUBNODE UNAVAILABLE (33)
- SUBNODE NOT FOUND IN CONFIG (34)
- ATTRIBUTE ERROR (35)
- CLASSPATH ERROR (36)
- CONNECTION FAILURE (37)
- FILTER SYNTAX ERROR (38)
- FILE NAME MISSING (39)
- SQL QUERY ERROR (40)
- SQL FILTER QUERY ERROR (41)
- SQL DB QUERY ERROR (42)
- SQL DB FILTER QUERY ERROR (43)
- PORT OPEN FAILED (44)
- ACCESS DENIED (45)

- TIMEOUT (46)
- NOT IMPLEMENTED (47)
- REQUESTED A BAD VALUE (48)
- RESPONSE TOO BIG (49)
- GENERAL RESPONSE ERROR (50)
- SCRIPT NONZERO RETURN (51)
- SCRIPT NOT FOUND (52)
- SCRIPT LAUNCH ERROR (53)
- CONF FILE DOES NOT EXIST (54)
- CONF FILE ACCESS DENIED (55)
- INVALID CONF FILE (56)
- EIF INITIALIZATION FAILED (57)
- CANNOT OPEN FORMAT FILE (58)
- FORMAT FILE SYNTAX ERROR (59)
- REMOTE HOST UNAVAILABLE (60)
- EVENT LOG DOES NOT EXIST (61)
- PING FILE DOES NOT EXIST (62)
- NO PING DEVICE FILES (63)
- PING DEVICE LIST FILE MISSING (64)
- SNMP MISSING PASSWORD (65)
- DISABLED (66)
- URLS FILE NOT FOUND (67)
- XML PARSE ERROR (68)
- NOT INITIALIZED (69)
- ICMP SOCKETS FAILED (70)
- DUPLICATE CONF FILE (71)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

ERROR_CODE or ERRCODE

Last Collection Start attribute

Description

The most recent time a data collection of this group started.

Type

Timestamp with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- NOT COLLECTED (0691231190000000)
- NOT COLLECTED (0000000000000001)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

LAST_COLLECTION_START or COLSTR

Last Collection Finished attribute

Description

The most recent time a data collection of this group finished.

Type

Timestamp with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- NOT COLLECTED (0691231190000000)
- NOT COLLECTED (0000000000000001)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

LAST_COLLECTION_FINISHED or COLFINI

Last Collection Duration attribute**Description**

The duration of the most recently completed data collection of this group in seconds.

Type

Real number (32-bit counter) with two decimal places of precision

Warehouse name

LAST_COLLECTION_DURATION or COLDURA

Average Collection Duration attribute**Description**

The average duration of all data collections of this group in seconds.

Type

Real number (32-bit counter) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- NO DATA (-100)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

AVERAGE_COLLECTION_DURATION or COLAVGD

Refresh Interval attribute**Description**

The interval at which this group is refreshed in seconds.

Type

Integer (32-bit counter)

Warehouse name

REFRESH_INTERVAL or REFRINT

Number of Collections attribute**Description**

The number of times this group has been collected since agent start.

Type

Integer (32-bit counter)

Warehouse name

NUMBER_OF_COLLECTIONS or NUMCOLL

Cache Hits attribute**Description**

The number of times an external data request for this group was satisfied from the cache.

Type

Integer (32-bit counter)

Warehouse name

CACHE_HITS or CACHEHT

Cache Misses attribute**Description**

The number of times an external data request for this group was not available in the cache.

Type

Integer (32-bit counter)

Warehouse name

CACHE_MISSES or CACHEMS

Cache Hit Percent attribute

Description

The percentage of external data requests for this group that were satisfied from the cache.

Type

Real number (32-bit counter) with two decimal places of precision

Warehouse name

CACHE_HIT_PERCENT or CACHPCT

Intervals Skipped attribute**Description**

The number of times a background data collection for this group was skipped because the previous collection was still running when the next one was due to start.

Type

Integer (32-bit counter)

Warehouse name

INTERVALS_SKIPPED or INTSKIP

Scheduler Parameters attribute group

The scheduler parameters attribute group contains the scheduler information for a virtual machine, including host name, virtual machine name, scheduler parameter name, scheduler parameter value, and scheduler parameter type.

Historical group

This attribute group is eligible for use with Tivoli Data Warehouse.

Attribute descriptions

The following list contains information about each attribute in the Scheduler Parameters attribute group:

Node attribute: This attribute is a key attribute.

Description

The managed system name of the agent.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

NODE

Timestamp attribute**Description**

The local time at the agent when the data was collected.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

TIMESTAMP

Host Name attribute: This attribute is a key attribute.

Description

The name of the host.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

HOST_NAME

Virtual Machine Name attribute: This attribute is a key attribute.**Description**

The name of the virtual machine.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VIRTUAL_MACHINE_NAME or VM_NAME

Scheduler Parameter Name attribute: This attribute is a key attribute.**Description**

The name of the scheduler parameter.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

SCHEDULER_PARAMETER_NAME or SCHP_NAME

Scheduler Parameter Value attribute**Description**

The value of the scheduler parameter as a string.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

SCHEDULER_PARAMETER_VALUE or SCHP_VALUE

Scheduler Parameter Type attribute**Description**

The original type of the value of the scheduler parameter, before conversion to a string, for example, integer, string, Boolean.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

SCHEDULER_PARAMETER_TYPE or SCHP_TYPE

Storage Pools attribute group

The storage pool attribute group contains information about storage pools, including storage pool name, containing host name, UUID, state, type, capacity, used, available, and percent used.

Historical group

This attribute group is eligible for use with Tivoli Data Warehouse.

Attribute descriptions

The following list contains information about each attribute in the Storage Pools attribute group:

Node attribute: This attribute is a key attribute.

Description

The managed system name of the agent.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

NODE

Timestamp attribute

Description

The local time at the agent when the data was collected.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

TIMESTAMP

Host Name attribute: This attribute is a key attribute.

Description

The name of the host.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

HOST_NAME

Storage Pool Name attribute: This attribute is a key attribute.

Description

The name of the storage pool.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

STORAGE_POOL_NAME or STGPL_NAME

Storage Pool UUID attribute

Description

The UUID of the storage pool.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise

Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

STORAGE_POOL_UUID or STGPL_UUID

Storage Pool State attribute

Description

The run state of the storage pool.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)
- Initializing (VIR_STORAGE_POOL_BUILDING)
- Degraded (VIR_STORAGE_POOL_DEGRADED)
- Inactive (VIR_STORAGE_POOL_INACTIVE)
- Running (VIR_STORAGE_POOL_RUNNING)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

STORAGE_POOL_STATE or STGPL_STAT

Storage Pool Type attribute

Description

The type of the storage pool, which is one of dir, fs, netfs, disk, iscsi, or logical.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

STORAGE_POOL_TYPE or STGPL_TYPE

Storage Pool Capacity (GB) attribute

Description

The storage capacity of the storage pool in GB.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

STORAGE_POOL_CAPACITY or STG_CAPCTY

Storage Pool Used (GB) attribute

Description

The amount of storage used from the storage pool in GB.

Type

Real number (64-bit gauge) with two decimal places of precision with

enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

STORAGE_POOL_USED or STG_USED

Storage Pool Available (GB) attribute

Description

The amount of storage available in the storage pool in GB.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

STORAGE_POOL_AVAILABLE or STG_AVAIL

Percent Used (%) attribute

Description

The amount of storage used from the storage pool as a percent of the total capacity.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

PERCENT_USED or PCT_USED

Virtual Machines attribute group

The virtual machine attribute group contains information about virtual machines, including virtual machine name, containing host name, UUID, state, number of CPUs, CPU time, memory used, life-cycle control, and CPU requirements.

Historical group

This attribute group is eligible for use with Tivoli Data Warehouse.

Attribute descriptions

The following list contains information about each attribute in the Virtual Machines attribute group:

Node attribute: This attribute is a key attribute.

Description

The managed system name of the agent.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

NODE

Timestamp attribute

Description

The local time at the agent when the data was collected.

Type

String

Source

The source for this attribute is the agent.

Warehouse name

TIMESTAMP

Host Name attribute: This attribute is a key attribute.

Description

The name of the host.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

HOST_NAME

Virtual Machine Name attribute: This attribute is a key attribute.

Description

The name of the virtual machine.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VIRTUAL_MACHINE_NAME or VM_NAME

Virtual Machine UUID attribute

Description

The UUID of the virtual machine.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VIRTUAL_MACHINE_UUID or VM_UUID

Virtual Machine State attribute

Description

The running state of the virtual machine.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)
- Blocked (VIR_DOMAIN_BLOCKED)

- Crashed (VIR_DOMAIN_CRASHED)
- No State (VIR_DOMAIN_NOSTATE)
- Paused (VIR_DOMAIN_PAUSED)
- Running (VIR_DOMAIN_RUNNING)
- Shutting Down (VIR_DOMAIN_SHUTDOWN)
- Shut Off (VIR_DOMAIN_SHUTOFF)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VIRTUAL_MACHINE_STATE or VM_STATE

Virtualization Type attribute

Description

The type of virtualization used by this virtual machine. hvm means full virtualization, which is used by KVM; linux means para virtualization, which is used by XEN.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

VIRTUALIZATION_TYPE or VIRT_TYPE

Number Of Virtual CPUs attribute

Description

The number of virtual CPUs in the virtual machine.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

NUMBER_OF_VIRTUAL_CPUS or VCPU_COUNT

Sample Timestamp attribute

Description

The date and time UTC of the most recent CPU measurement.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (0000-00-00 00:00:00)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

SAMPLE_TIMESTAMP or SAMP_TS

Sample Time Delta (sec) attribute

Description

The time duration of the most recent measurement period in seconds.

Type

Real number (64-bit gauge) with two decimal places of precision with

enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

SAMPLE_TIME_DELTA or SAMP_DELTA

Total CPU Time (sec) attribute

Description

The CPU time used in seconds.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_TIME

CPU Time Delta (sec) attribute

Description

The CPU time used during the most recent measurement period in seconds.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_TIME_DELTA or CPU_DELTA

CPUs Consumed attribute

Description

The equivalent number of real CPUs consumed by this virtual machine.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPUS_CONSUMED or CPU_CNSMD

CPU Percent attribute

Description

The composite CPU percent of this virtual machine across all virtual CPUs.

Type

Real number (64-bit gauge) with two decimal places of precision with

enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_PERCENT or CPU_PCT

Memory Allocated (GB) attribute

Description

The maximum memory allocated and available for the virtual machine in GB.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

MEMORY_ALLOCATED or MEMRY_ALLO

Memory Used (GB) attribute

Description

The memory used by the virtual machine in GB.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

MEMORY_USED or MEMRY_USED

Memory Percent attribute

Description

The percent of allocated memory that is used by this virtual machine.

Type

Real number (64-bit gauge) with two decimal places of precision with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

MEMORY_PERCENT or MEMRY_PCT

Action On PowerOff attribute

Description

The action to take when the virtual machine requests a power-off.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

ACTION_ON_POWEROFF or ACTN_PWROF

Action On Reboot attribute

Description

The action to take when the virtual machine requests a restart.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

ACTION_ON_REBOOT or ACTN_RBOOT

Action On Crash attribute

Description

The action to take when the virtual machine crashes.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

ACTION_ON_CRASH or ACTN_CRASH

CPU Match attribute

Description

How strictly the virtual CPU provided to the virtual machine must match these requirements.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)
- Minimum (minimum)
- Exact (exact)
- Strict (strict)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_MATCH

CPU Model attribute

Description

The CPU model requested for the virtual machine.

Type

String with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (unavailable)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_MODEL

Number of Sockets attribute

Description

The requested number of CPU sockets for the virtual machine.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

NUMBER_OF_SOCKETS or SKTS_COUNT

Cores Per Socket attribute

Description

The requested number of cores per socket for the virtual machine.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CORES_PER_SOCKET or CORES_PER

Threads Per Core attribute

Description

The requested number of threads per core for the virtual machine.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

THREADS_PER_CORE or THRDS_PER

CPU Shares attribute

Description

The scheduler parameter, `cpu_shares`, for the virtual machine.

Type

Integer (64-bit gauge) with enumerated values. The strings are displayed in the Tivoli Enterprise Portal. The warehouse and queries return the values that are shown in parentheses. The following values are defined:

- Unavailable (-1)

Any other value is the value that is returned by the agent in the Tivoli Enterprise Portal.

Warehouse name

CPU_SHARES

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 by the historical data in each attribute group that is collecting historical data. Required disk storage is an important factor when you are defining data collection rules and your strategy for historical data collection.

The Capacity planning for historical data table provides the following information, which is required to calculate disk space for this monitoring agent:

Table Table name as it is displayed in the warehouse database, if the attribute group is configured to be written to the warehouse. The table name listed here corresponds to the table name in “Attribute groups for the monitoring agent” on page 21.

Attribute group

Name of the attribute group that is used to create the table in the warehouse database if it is short enough to fit in the table naming constraints of the database that is being used for the warehouse. The attribute group name listed here corresponds to the Warehouse table name in “Attribute groups for the monitoring agent” on page 21.

Bytes per row (agent)

Estimate of the record length for each row or instance that is written to the agent disk for historical data collection. This estimate can be used for agent disk space planning purposes.

Database bytes per row (warehouse)

Estimate of the record length for detailed records that are written to the warehouse database, if the attribute group is configured to be written to the warehouse. Detailed records are records 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 row (warehouse)

Estimate of the record length for aggregate records that are 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.

In addition to the information in the tables, you must know the number of rows of data that you plan to collect. An attribute group can have single or multiple rows of data, depending on the application environment that is being monitored. For example, if your attribute group monitors each processor in your computer and you have a dual processor computer, the number of rows is two.

Table 1. Capacity planning for historical data logged by the Linux Kernel-based Virtual Machines agent

Table	Attribute group	Bytes per row (agent)	Database bytes per row (warehouse)	Aggregate bytes per row (warehouse)
KV1HOSTAG	KV1_HOSTS	810	845	1336
KV1HOSTCG	KV1_HOST_CPU	459	568	1389
KV1HOSTMG	KV1_HOST_MEMORY	443	558	1301
KV1POBJST	KV1_PERFORMANCE_OBJECT_STATUS	352	399	664
KV1STGPLAG	KV1_STORAGE_POOLS	730	767	1008
KV1VMACHAG	KV1_VIRTUAL_MACHINES	906	989	1606

For more information about historical data collection, see *Managing historical data* in the *IBM Tivoli Monitoring Administrator's Guide*.

Chapter 5. Situations reference

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 or from the command-line interface using the `tacmd` commands for situations. You can manage private situations in the private configuration XML file.

About situations

The monitoring agents that you use to monitor your system environment include a set of predefined situations that you can use as-is. You can also 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 IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines. You can change the conditions or values being monitored by a predefined situation to the conditions or values best suited to your enterprise.

You can display predefined situations and create your own situations using the Situation editor. The Situation editor initially lists the situations associated with the Navigator item that you selected. When you click a situation name or create a situation, a panel opens with the following tabs:

Formula

Formula describing the condition being tested.

Distribution

List of managed systems (operating systems, subsystems, or applications) to which the situation can be distributed. All the Linux Kernel-based Virtual Machines agent managed systems are assigned by default.

Expert advice

Comments and instructions to be read in the event workspace.

Action

Command to be sent to the system.

EIF Customize forwarding of the event to an Event Integration Facility receiver. (Available when the Tivoli Enterprise Monitoring Server is configured to forward events.)

Until Options to close the event after a period of time, or when another situation becomes true.

Additional information about situations

The *Tivoli Enterprise Portal 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 and information about each individual situation for this monitoring agent, see "Predefined situations."

Predefined situations

The monitoring agent contains predefined situations, which are organized by Navigator item.

- Linux Kernel-based Virtual Machines
 - Not applicable
- Host

- KV1_Host_CPU_Over_Commit_Info
- KV1_Host_CPU_Over_Commit_Warn
- KV1_Host_CPU_Over_Commit_Crit
- KV1_Host_CPU_Pct_High_Warn
- KV1_Host_CPU_Pct_High_Crit
- KV1_Host_Mem_Pct_High_Warn
- KV1_Host_Mem_Pct_High_Crit
- Storage Pool
 - Not applicable
- Virtual Machine
 - KV1_VM_CPU_Pct_High_Warn
 - KV1_VM_CPU_Pct_High_Crit
 - KV1_VM_Mem_Pct_High_Warn
 - KV1_VM_Mem_Pct_High_Crit

Situation descriptions

Each situation description provides information about the situation that you can use to monitor the condition of systems in your network.

The situation descriptions provide the following information:

Description

Information about the conditions that the situation tests.

Formula

Syntax that contains one or more logical expressions that describe the conditions for the situation to monitor.

Distribution

Whether the situation is automatically distributed to instances of the agent or is available for manual distribution.

Run at startup

Whether the situation starts monitoring when the agent starts.

Sampling interval

Number of seconds that elapse between one sample of data that the monitoring agent collects for the server and the next sample.

Situation persistence

Whether the conditions specified in the situation evaluate to "true" for the defined number of occurrences in a row before the situation is raised. The default of one means that no persistence-checking takes place.

Severity

Severity of the predefined events: Warning, Informational, or Critical.

Clearing conditions

Controls when a true situation closes: after a period, when another situation is true, or whichever occurs first if both are selected.

Linux Kernel-based Virtual Machines Navigator item

No predefined situations are included for this Navigator item.

Host Navigator item

The situation descriptions are organized by the Navigator item to which the situations are relevant.

KV1_Host_CPU_Over_Commit_Info situation

Description

The CPU of your host is over committed.

The situation is evaluated for each distinct value of the HOST_NAME attribute.

Formula

```
*IF *VALUE KV1_HOST_CPU.Percent_CPU_Allocated *GE 100 *AND *VALUE  
KV1_HOST_CPU.Percent_CPU_Allocated *LT 125
```

See “Attributes in each attribute group” on page 22 for descriptions of the attributes in this formula.

Distribution

This situation is available for distribution.

Run at startup

No

Sampling interval

15 minutes

Situation persistence

The number of times the conditions of the situation must occur for the situation to be true is 1.

Error conditions

Informational

Clearing conditions

The situation clears when the condition becomes false.

KV1_Host_CPU_Over_Commit_Warn situation

Description

The CPU of your host is moderately over committed.

The situation is evaluated for each distinct value of the HOST_NAME attribute.

Formula

```
*IF *VALUE KV1_HOST_CPU.Percent_CPU_Allocated *GE 125 *AND *VALUE  
KV1_HOST_CPU.Percent_CPU_Allocated *LT 150
```

See “Attributes in each attribute group” on page 22 for descriptions of the attributes in this formula.

Distribution

This situation is available for distribution.

Run at startup

No

Sampling interval

15 minutes

Situation persistence

The number of times the conditions of the situation must occur for the situation to be true is 1.

Error conditions

Warning

Clearing conditions

The situation clears when the condition becomes false.

KV1_Host_CPU_Over_Commit_Crit situation

Description

The CPU of your host is highly over committed.

The situation is evaluated for each distinct value of the HOST_NAME attribute.

Formula

```
*IF *VALUE KV1_HOST_CPU.Percent_CPU_Allocated *GE 150
```

See “Attributes in each attribute group” on page 22 for descriptions of the attributes in this formula.

Distribution

This situation is available for distribution.

Run at startup

No

Sampling interval

15 minutes

Situation persistence

The number of times the conditions of the situation must occur for the situation to be true is 1.

Error conditions

Critical

Clearing conditions

The situation clears when the condition becomes false.

KV1_Host_CPU_Pct_High_Warn situation

Description

Your host has chronically low available CPU capacity.

The situation is evaluated for each distinct value of the HOST_NAME attribute.

Formula

```
*IF *VALUE KV1_HOST_CPU.CPU_Percent_Consumed *GE 85 *UNTIL ( *SIT
KV1_Host_CPU_Pct_High_Crit )
```

See “Attributes in each attribute group” on page 22 for descriptions of the attributes in this formula.

Distribution

This situation is available for distribution.

Run at startup

No

Sampling interval

15 minutes

Situation persistence

The number of times the conditions of the situation must occur for the situation to be true is 3.

Error conditions

Warning

Clearing conditions

The situation clears when the condition becomes false.

KV1_Host_CPU_Pct_High_Crit situation

Description

Your host has persistently low available CPU capacity.

The situation is evaluated for each distinct value of the HOST_NAME attribute.

Formula

```
*IF *VALUE KV1_HOST_CPU.CPU_Percent_Consumed *GE 85
```

See “Attributes in each attribute group” on page 22 for descriptions of the attributes in this formula.

Distribution

This situation is available for distribution.

Run at startup

No

Sampling interval

15 minutes

Situation persistence

The number of times the conditions of the situation must occur for the situation to be true is 4.

Error conditions

Critical

Clearing conditions

The situation clears when the condition becomes false.

KV1_Host_Mem_Pct_High_Warn situation**Description**

The memory utilization of the host is persistently high.

The situation is evaluated for each distinct value of the HOST_NAME attribute.

Formula

```
*IF *VALUE KV1_HOST_MEMORY.Percent_Memory_Used *GE 95 *UNTIL ( *SIT
KV1_Host_Mem_Pct_High_Crit )
```

See “Attributes in each attribute group” on page 22 for descriptions of the attributes in this formula.

Distribution

This situation is available for distribution.

Run at startup

No

Sampling interval

15 minutes

Situation persistence

The number of times the conditions of the situation must occur for the situation to be true is 2.

Error conditions

Warning

Clearing conditions

The situation clears when the condition becomes false.

KV1_Host_Mem_Pct_High_Crit situation**Description**

The memory utilization of the host is chronically high.

The situation is evaluated for each distinct value of the HOST_NAME attribute.

Formula

```
*IF *VALUE KV1_HOST_MEMORY.Percent_Memory_Used *GE 95
```

See “Attributes in each attribute group” on page 22 for descriptions of the attributes in this formula.

Distribution

This situation is available for distribution.

Run at startup

No

Sampling interval

15 minutes

Situation persistence

The number of times the conditions of the situation must occur for the situation to be true is 3.

Error conditions

Critical

Clearing conditions

The situation clears when the condition becomes false.

Storage Pool Navigator item

No predefined situations are included for this Navigator item.

Virtual Machine Navigator item

The situation descriptions are organized by the Navigator item to which the situations are relevant.

KV1_VM_CPU_Pct_High_Warn situation

Description

The CPU utilization of the virtual machine is persistently high.

The situation is evaluated for each distinct value of the HOST_NAME attribute.

Formula

```
*IF *VALUE KV1_VIRTUAL_MACHINES.CPU_Percent *GE 95 *UNTIL ( *SIT  
KV1_VM_CPU_Pct_High_Crit )
```

See “Attributes in each attribute group” on page 22 for descriptions of the attributes in this formula.

Distribution

This situation is available for distribution.

Run at startup

No

Sampling interval

15 minutes

Situation persistence

The number of times the conditions of the situation must occur for the situation to be true is 2.

Error conditions

Warning

Clearing conditions

The situation clears when the condition becomes false.

KV1_VM_CPU_Pct_High_Crit situation

Description

The CPU utilization of the virtual machine is chronically high.

The situation is evaluated for each distinct value of the HOST_NAME attribute.

Formula

```
*IF *VALUE KV1_VIRTUAL_MACHINES.CPU_Percent *GE 95
```

See “Attributes in each attribute group” on page 22 for descriptions of the attributes in this formula.

Distribution

This situation is available for distribution.

Run at startup

No

Sampling interval

15 minutes

Situation persistence

The number of times the conditions of the situation must occur for the situation to be true is 3.

Error conditions

Critical

Clearing conditions

The situation clears when the condition becomes false.

KV1_VM_Mem_Pct_High_Warn situation

Description

The memory utilization of the VM is persistently high.

The situation is evaluated for each distinct value of the HOST_NAME attribute.

Formula

```
*IF *VALUE KV1_VIRTUAL_MACHINES.Memory_Percent *GE 95 *UNTIL ( *SIT  
KV1_VM_Mem_Pct_High_Crit )
```

See “Attributes in each attribute group” on page 22 for descriptions of the attributes in this formula.

Distribution

This situation is available for distribution.

Run at startup

No

Sampling interval

15 minutes

Situation persistence

The number of times the conditions of the situation must occur for the situation to be true is 2.

Error conditions

Warning

Clearing conditions

The situation clears when the condition becomes false.

KV1_VM_Mem_Pct_High_Crit situation

Description

The memory utilization of the VM is chronically high.

The situation is evaluated for each distinct value of the HOST_NAME attribute.

Formula

*IF *VALUE KV1_VIRTUAL_MACHINES.Memory_Percent *GE 95

See “Attributes in each attribute group” on page 22 for descriptions of the attributes in this formula.

Distribution

This situation is available for distribution.

Run at startup

No

Sampling interval

15 minutes

Situation persistence

The number of times the conditions of the situation must occur for the situation to be true is 3.

Error conditions

Critical

Clearing conditions

The situation clears when the condition becomes false.

Chapter 6. Take Action commands reference

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

About Take Action commands

When included in a situation, the command runs when the situation becomes true. A Take Action command in a situation is also referred to as *reflex automation*. When you enable a Take Action command in a situation, you automate a response to system conditions. For example, you can use a Take Action command to send a command to restart a process on the managed system or to send a text message to a cell phone.

In advanced automation, policies are used to take 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, the Tivoli Enterprise Portal receives return-code feedback, and advanced automation logic responds with subsequent activities that are prescribed by the feedback.

A basic Take Action command shows the return code of the operation in a message box that is displayed after the action is completed or in a log file. After you close this window, no further information is available for this action.

Additional information about Take Action commands

For more information about working with Take Action commands, see *Take Action commands* in the *Tivoli Enterprise Portal User's Guide*.

Predefined Take Action commands

Not all agents have predefined Take Action commands. But you can create Take Action commands for any agent.

The IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines does not provide predefined Take Action commands.

Chapter 7. Policies reference

Policies are used as an advanced automation technique for implementing more complex workflow strategies than you can create through simple automation. All agents do not provide predefined policies, but you can create policies for any agent.

A *policy* is a set of automated system processes that can take 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, the Tivoli Enterprise Portal receives return-code feedback, and advanced automation logic responds with subsequent activities prescribed by the feedback.

For more information about working with policies, see *Automation with policies* in the *Tivoli Enterprise Portal 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.

Predefined policies

Not all agents have predefined policies. But you can create policies for any agent.

The IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines does not provide predefined policies.

Chapter 8. Troubleshooting

Problems can be related to IBM Tivoli Monitoring or the specific agent that you are using.

For general troubleshooting information, see the *IBM Tivoli Monitoring Troubleshooting Guide*. For other problem-solving options, see “Support information” on page 88.

You can resolve some problems by ensuring that your system matches the system requirements listed in the Prerequisites topic for the agent in the information center, or in the Requirements topic of the agent user's guide.

The following activities can help you find a solution to the problem you are having:

- “Gathering product information for IBM Software Support”
- “Using logging” on page 66
- “Consulting the lists of identified problems and workarounds” on page 66

Gathering product information for IBM Software Support

Before contacting IBM Software Support about a problem you are experiencing with this product, gather the information shown in Table 2.

Table 2. 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 “Principal trace log files” on page 67 for lists of all trace log files and their locations. For general information about the IBM Tivoli Monitoring environment, see the <i>Tivoli Enterprise Portal User's Guide</i> .
Linux Kernel-based Virtual Machines information	Version number and patch level
Operating system	Operating system version number and patch level
Messages	Messages and other information displayed on the screen
Version numbers for IBM Tivoli Monitoring	Version number of the following members of the monitoring environment: <ul style="list-style-type: none">• IBM Tivoli Monitoring. Also provide the patch level, if available.• IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines
Screen captures	Screen captures of incorrect output, if any
(UNIX systems only) Core dump files	If the system stops on UNIX systems, collect the core dump file from the <i>install_dir/bin</i> directory, where <i>install_dir</i> is the directory where you installed the monitoring agent.

You can use the `pdcollect` tool to collect the most commonly used information from a system. This tool gathers log files, configuration information, version information, and other data. For more information about using this tool, see the “`pdcollect` tool” in the *IBM Tivoli Monitoring Troubleshooting Guide*.

For information about working with IBM Software Support, see IBM Support Portal Service Requests and PMRs ([http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_\(general\)](http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_(general))).

Using logging

Logging is the primary troubleshooting feature in the Linux Kernel-based Virtual Machines agent. *Logging* refers to the text messages and trace data that is generated by the Linux Kernel-based Virtual Machines 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.

Consulting the lists of identified problems and workarounds

Known problems are organized into types such as those in the following list to make them easier to locate:

- Installation and configuration
- General usage and operation
- Display of monitoring data
- Take Action commands

Information about symptoms and detailed workarounds for these types of problems is located in “Problems and workarounds” on page 77.

For general troubleshooting information, see the *IBM Tivoli Monitoring Troubleshooting Guide*.

Trace logging

Trace logs are used to capture information about the operating environment when component software fails to operate as designed.

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 information to learn how to configure and use trace logging:

- “Principal trace log files” on page 67
- “Examples: Using trace logs” on page 70
- “Setting RAS trace parameters by using the GUI” on page 71

Note: The documentation refers to the RAS facility in IBM Tivoli Monitoring as “RAS1.”

IBM Software Support personnel use the information captured by trace logging to trace a problem to its source or to determine why an error occurred. All components in the IBM Tivoli Monitoring environment have a default tracing level. The tracing level can be changed on a per-component level to adjust the type of trace information collected, the degree of trace detail, the number of trace logs to be kept, and the amount of disk space used for tracing.

Overview of log file management

Knowing the naming conventions for log files helps you to find the files.

Agent log file naming conventions

Table 3 provides the names, locations, and descriptions of IBM Tivoli Monitoring general RAS1 log files. The log file names for the Linux Kernel-based Virtual Machines agent adhere to the following naming convention:

Windows systems

hostname_productcode_instance-name_program_HEXtimestamp-nn.log

Linux and UNIX systems

hostname_productcode_instance-name_program_HEXtimestamp-nn.log

Where:

hostname

Host name of the computer where the monitoring component is running.

productcode

Two-character product code. For IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines, the product code is v1.

instance-name

Instance name of the agent.

program

Name of the program being run.

HEXtimestamp

Hexadecimal time stamp representing the time at which the program started.

nn

Rolling log suffix.

Principal trace log files

Trace log files are located on various systems.

Table 3 contains locations, file names, and descriptions of trace logs that can help determine the source of problems with agents.

Table 3. Trace log files for troubleshooting agents

System where log is located	File name and path	Description
On the Tivoli Enterprise Monitoring Server	<ul style="list-style-type: none"> • Windows: The file IBM Tivoli Monitoring <i>timestamp.log</i> in the <i>install_dir\InstallITM</i> path • UNIX: The <i>candle_installation.log</i> file in the <i>install_dir/logs</i> path • Linux: The <i>candle_installation.log</i> file in the <i>install_dir/logs</i> 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.
On the Tivoli Enterprise Monitoring Server	The <i>Warehouse_Configuration.log</i> file is in the following location on Windows systems: <i>install_dir\InstallITM</i>	Provides details about the configuration of data warehousing for historical reporting.

Table 3. Trace log files for troubleshooting agents (continued)

System where log is located	File name and path	Description
On the Tivoli Enterprise Monitoring Server	<p>The name of the RAS log file is as follows:</p> <ul style="list-style-type: none"> • Windows: <i>install_dir\logs\hostname_ms_timestamp-<i>nn</i>.log</i> • UNIX: <i>install_dir/logs/hostname_ms_timestamp-<i>nn</i>.log</i> • Linux: <i>install_dir/logs/hostname_ms_timestamp-<i>nn</i>.log</i> <p>Note: File names for RAS1 logs include a hexadecimal time stamp.</p> <p>Also on UNIX systems, a log with a decimal time stamp is provided: <i>hostname_productcode_timestamp.log</i> and <i>hostname_productcode_timestamp.pid nnnnn</i> in the <i>install_dir/logs</i> path, where <i>nnnnn</i> is the process ID number.</p>	Traces activity on the monitoring server.
On the Tivoli Enterprise Portal Server	<p>The name of the RAS log file is as follows:</p> <ul style="list-style-type: none"> • Windows: <i>install_dir\logs\hostname_cq_HEXtimestamp-<i>nn</i>.log</i> • UNIX: <i>install_dir/logs/hostname_cq_HEXtimestamp-<i>nn</i>.log</i> • Linux: <i>install_dir/logs/hostname_cq_HEXtimestamp-<i>nn</i>.log</i> <p>Note: File names for RAS1 logs include a hexadecimal time stamp.</p> <p>Also on UNIX systems, a log with a decimal time stamp is provided: <i>hostname_productcode_timestamp.log</i> and <i>hostname_productcode_timestamp.pidnnnn</i> in the <i>install_dir/logs</i> path, where <i>nnnnn</i> is the process ID number.</p>	Traces activity on the portal server.
On the Tivoli Enterprise Portal Server	<p>The <i>teps_odbc.log</i> file is located in the following path:</p> <ul style="list-style-type: none"> • Windows: <i>install_dir\Install\ITM</i> • UNIX: <i>install_dir/logs</i> • Linux: <i>install_dir/logs</i> 	When you enable historical reporting, this log file traces the status of the warehouse proxy agent.

Table 3. Trace log files for troubleshooting agents (continued)

System where log is located	File name and path	Description
On the computer that hosts the monitoring agent	<p>The RAS1 log files are as follows:</p> <ul style="list-style-type: none"> • UNIX: <i>hostname_v1_instance_name_kv1agent_HEXtimestamp-nn.log</i> in the <i>install_dir/logs</i> directory • Linux: <i>hostname_v1_instance_name_kv1agent_HEXtimestamp-nn.log</i> in the <i>install_dir/logs</i> directory <p>These logs are in the following directories:</p> <ul style="list-style-type: none"> • UNIX: <i>install_dir/logs</i> • Linux: <i>install_dir/logs</i> <p>On Linux systems, the following additional logs are provided:</p> <ul style="list-style-type: none"> – <i>hostname_v1_timestamp.log</i> – <i>hostname_v1_timestamp.pidnnnn</i> in the <i>install_dir/logs</i> path, where <i>nnnnn</i> is the process ID number 	Traces activity of the monitoring agent.
On the computer that hosts the monitoring agent	<p>The agent operations log files are as follows:</p> <p><i>instance_hostnameV1.LG0</i> is the current log created when the agent was started.</p> <p><i>instance_hostname_V1.LG1</i> is the backup of the previous log.</p> <p>These logs are in the following directory depending on the operating system that you are using:</p> <ul style="list-style-type: none"> • Linux: <i>install_dir/logs</i> • UNIX: <i>install_dir/logs</i> 	<p>Shows whether the agent could 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.</p> <p>IBM Tivoli Monitoring generates one backup copy of the *.LG0 file with the tag .LG1. View the .LG1 tag to learn the following details regarding the <i>previous</i> monitoring session:</p> <ul style="list-style-type: none"> • Status of connectivity with the monitoring server • Situations that were running • The success or failure status of Take Action commands
On the computer that hosts the monitoring agent	<p>The Take Action command log files are as follows:</p> <ul style="list-style-type: none"> • <i>host_v1_instance_takeactioncommand .log</i> <p>The logs are in the following directories:</p> <ul style="list-style-type: none"> • UNIX: <i>install_dir /logs</i> • Linux: <i>install_dir /logs</i> 	Traces activity each time a Take Action command runs. For example, when a hypothetical start_command Take Action command runs, IBM Tivoli Monitoring generates a <i>start_command.log</i> file.

Table 3. Trace log files for troubleshooting agents (continued)

System where log is located	File name and path	Description
On the computer that hosts the monitoring agent	<p>The Take Action command log files are as follows:</p> <ul style="list-style-type: none"> • <code>kv1_data_provider_actions_instance_n.log</code> <p>The logs are in the following directories:</p> <ul style="list-style-type: none"> • UNIX: <code>install_dir/logs</code> • Linux: <code>install_dir/logs</code> 	Traces activity each time a Take Action command runs. All predefined Take Action commands are logged into this file.
<p>Definitions of variables:</p> <ul style="list-style-type: none"> • <i>timestamp</i> is a time stamp with a format that includes year (y), month (m), day (d), hour (h), and minute (m), as follows: yyyymmdd hhmm • <i>HEXtimestamp</i> is a hexadecimal representation of the time at which the process was started. • <i>install_dir</i> represents the directory path where you installed the IBM Tivoli Monitoring component. <i>install_dir</i> can represent a path on the computer that hosts the monitoring system, the monitoring agent, or the portal. • <i>instance</i> refers to the name of the database instance that you are monitoring. • <i>instance_name</i> refers to the name of the agent instance. • <i>hostname</i> refers to the name of the computer on which the IBM Tivoli Monitoring component runs. • <i>nn</i> represents the circular sequence in which logs are rotated. this value includes a range from 1 - 5, by default. The first is always retained because it includes configuration parameters. • <i>productcode</i> specifies the product code, for example, um for Universal Agent or nt for Windows systems. 		

For more information about the complete set of trace logs that are maintained on the monitoring server, see the *IBM Tivoli Monitoring Installation and Setup Guide*.

Examples: Using trace logs

You can open trace logs in a text editor to learn some basic facts about your IBM Tivoli Monitoring environment.

IBM Software Support applies specialized knowledge to analyze trace logs to determine the source of problems. The following examples are from the Tivoli Enterprise Monitoring Server log.

Example one

This excerpt shows the typical log for a failed connection between a monitoring agent and a monitoring server with the host name **server1a**:

```
(Thursday, August 11, 2005, 08:21:30-{94C}kdc10c1.c,105,"KDCL0_ClientLookup") status=1c020006,
"location server unavailable", ncs/KDC1_STC_SERVER_UNAVAILABLE
(Thursday, August 11, 2005, 08:21:35-{94C}kraarreg.cpp,1157,"LookupProxy") Unable to connect to
broker at ip.pipe:: status=0, "success", ncs/KDC1_STC_OK
(Thursday, August 11, 2005, 08:21:35-{94C}kraarreg.cpp,1402,"FindProxyUsingLocalLookup") Unable
to find running CMS on CT_CMSLIST <IP.PIPE:#server1a>
```

Example two

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

```
(42C039F9.0000-6A4:kpxreqhb.cpp,649,"HeartbeatInserter") Remote node SERVER5B:V1 is ON-LINE.
.
.
(42C3079B.0000-6A4:kpxreqhb.cpp,644,"HeartbeatInserter") Remote node SERVER5B:V1 is OFF-LINE.
```

See the following key points about the preceding excerpts:

- The monitoring server appends the **V1** product code to the server name to form a unique name (SERVER5B:V1) for this instance of the IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines. By using this unique name, you can 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 by using the GUI” provide these entries.

On Windows systems, you can use the following alternate method to view trace logs:

1. In the Windows **Start** menu, click **Program Files > IBM Tivoli Monitoring > Manage Tivoli Enterprise Monitoring Services**. The Manage Tivoli Enterprise Monitoring Services window is displayed.
2. Right-click a component and click **Advanced > View Trace Log** in the menu. For example, if you want to view the trace log of the IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines, right-click the name of that agent in the window. You can also use the viewer to access remote logs.

Note: The viewer converts time stamps in the logs to a format that is easier to read.

RAS trace parameters

Pinpoint a problem by setting detailed tracing of individual components of the monitoring agent and modules

See “Overview of log file management” on page 66 to ensure that you understand log rolling and can reference the correct log files when you manage log file generation.

Setting RAS trace parameters by using the GUI

On Windows systems, you can use the graphical user interface to set trace options.

About this task

The IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines uses RAS1 tracing and generates the logs described in Table 3 on page 67. The default RAS1 trace level is ERROR.

Procedure

1. Open the Manage Tivoli Enterprise Monitoring Services window.
2. Select **Advanced > Edit Trace Parm**s. The Tivoli Enterprise Monitoring Server Trace Parameters window is displayed.
3. Select a new trace setting in the pull-down menu in the **Enter RAS1 Filters** field or type a valid string.
 - General error tracing. KBB_RAS1=ERROR
 - Intensive error tracing. KBB_RAS1=ERROR (UNIT:kv1 ALL)
 - Maximum error tracing. KBB_RAS1=ERROR (UNIT:kv1 ALL) (UNIT:kra ALL)

Note: As this example shows, you can set multiple RAS tracing options in a single statement.

4. Modify the value for Maximum Log Size Per File (MB) to change the log file size (changes LIMIT value).
5. Modify the value for Maximum Number of Log Files Per Session to change the number of log files per startup of a program (changes COUNT value).
6. Modify the value for Maximum Number of Log Files Total to change the number of log files for all startups of a program (changes MAXFILES value).
7. Optional: Click Y (Yes) in the **KDC_DEBUG Setting** menu to log information that can help you diagnose communications and connectivity problems between the monitoring agent and the monitoring server. The **KDC_DEBUG** setting and the **Maximum error tracing** setting can generate a large amount of trace logging. Use these settings only temporarily, while you are troubleshooting problems. Otherwise, the logs can occupy excessive amounts of hard disk space.
8. Click **OK**. You see a message reporting a restart of the monitoring agent so that your changes take effect.

What to do next

Monitor the size of the logs directory. Default behavior can generate a total of 45 - 60 MB for each agent that is running on a computer. For example, each database instance that you monitor can generate 45 - 60 MB of log data. See the "Procedure" section to learn how to adjust file size and numbers of log files to prevent logging activity from occupying too much disk space.

Regularly prune log files other than the RAS1 log files in the logs directory. Unlike the RAS1 log files that are pruned automatically, other log types can grow indefinitely, for example, the logs in Table 3 on page 67 that include a process ID number (PID).

Use collector trace logs as an additional source of troubleshooting information.

Note: The **KDC_DEBUG** setting and the **Maximum error tracing** setting can generate a large amount of trace logging. Use these settings only temporarily while you are troubleshooting problems. Otherwise, the logs can occupy excessive amounts of hard disk space.

Manually setting RAS trace parameters

You can manually edit the RAS1 trace logging parameters.

About this task

The Linux Kernel-based Virtual Machines agent uses RAS1 tracing and generates the logs described in Table 3 on page 67. The default RAS1 trace level is ERROR.

Procedure

1. Open the trace options file:
 - **UNIX systems:**

```
install_dir /config/v1_instance_name.config
```
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: **KBB_RAS1=ERROR (UNIT:kv1 ALL) (UNIT:kra ALL)** See for detailed information.
3. Edit the line that begins with **KBB_RAS1_LOG=** to manage the generation of log files:
 - **MAXFILES:** The total number of files that are to be kept for all startups of a specific program. When this value is exceeded, the oldest log files are discarded. The default value is 9.
 - **LIMIT:** The maximum size, in megabytes (MB) of a RAS1 log file. The default value is 5.
 - IBM Software Support might guide you to modify the following parameters:
 - **COUNT:** The number of log files to keep in the rolling cycle of one program startup. The default is 3.

- **PRESERVE:** The number of files that are not to be reused in the rolling cycle of one program startup. The default value is 1.

Note: The **KBB_RAS1_LOG** parameter also provides for the specification of the log file directory, log file name, and the inventory control file directory and name. Do not modify these values or log information can be lost.

4. Restart the monitoring agent so that your changes take effect.

What to do next

Monitor the size of the logs directory. Default behavior can generate a total of 45 - 60 MB for each agent that is running on a computer. For example, each database instance that you monitor can generate 45 - 60 MB of log data. See the "Procedure" section to learn how to adjust file size and numbers of log files to prevent logging activity from occupying too much disk space.

Regularly prune log files other than the RAS1 log files in the logs directory. Unlike the RAS1 log files that are pruned automatically, other log types can grow indefinitely, for example, the logs in Table 3 on page 67 that include a process ID number (PID).

Use collector trace logs as an additional source of troubleshooting information.

Note: The **KDC_DEBUG** setting and the **Maximum error tracing** setting can generate a large amount of trace logging. Use these settings only temporarily while you are troubleshooting problems. Otherwise, the logs can occupy excessive amounts of hard disk space.

Dynamic modification of trace settings

You can dynamically modify the trace settings for an IBM Tivoli Monitoring component, such as, Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, most monitoring agents, and other components. You can access these components, except for a few monitoring agents, from the tracing utility.

Dynamic modification of the trace settings is the most efficient method, because you can do it without restarting the component. Settings take effect immediately. Modifications by this method are not persistent.

Note: When the component is restarted, the trace settings are read again from the .env file. Dynamically modifying these settings does not change the settings in the .env files. To modify these trace settings permanently, modify them in the .env files.

ras1

Run this command to modify the trace settings for a Tivoli Monitoring component.

The syntax is as follows:

```
ras1 set|list (UNIT|COMP: class_name ANY|ALL|Detail|ERROR|Flow|INPUT|Metrics|OUTPUT|STATE)
{(UNIT|COMP: class_name ANY|ALL|Detail|ERROR|Flow|INPUT|Metrics|OUTPUT|STATE)}
```

You can specify more than one component class to which to apply the trace settings.

Command options

set

Turns on or off tracing depending upon the value of its parameters. If the parameter is **ANY**, it turns it off. All other parameters turn on tracing based on the specified type or level.

list

Displays the default level and type of tracing that is set by default.

Parameters

The parameters that determine the component classes to which to apply the trace settings are as follows:

COMP: *class_name*

Modifies the trace setting for the name of the component class, as specified by *class_name* , for example, COMP:KDH. The output contains trace for the specified class.

UNIT: *class_name*

Modifies the trace setting for any unit that starts with the specified *class_name* value, for example, UNIT: kra. The output contains trace for any unit that begins with the specified filter pattern.

The parameters that determine the trace level and type are as follows:

ALL

Displays all trace levels, including every trace point defined for the component. This setting might result in a large amount of trace, so specify other parameters to exclude unwanted trace. You might require the **ALL** parameter to isolate a problem, which is the equivalent to setting "Error Detail Flow State Input Output Metrics".

ANY

Turns off tracing.

Detail

Displays detailed information about each function.

When entered with the *list* option, the trace is tagged with Det.

ERROR

Logs internal error conditions.

When entered with the *list* option, the trace is tagged with ER. The output can also be tagged with EVERYE+EVERYU+ER.

Flow

Displays control flow data for each function entry and exit.

When entered with the *list* option, the trace is tagged with Fl.

INPUT

Displays input data for each function.

When entered with the *list* option, the trace is tagged with IN.

Metrics

Displays metrics on each function.

When entered with the *list* option, the trace is tagged with ME.

OUTPUT

Displays output data for each function.

When entered with the *list* option, the trace is tagged with OUT.

State

Displays the status for each function.

When entered with the *list* option, the trace is tagged with St.

Example

If you enter `ras1 set (COMP:KDH ALL) (COMP:ACF1 ALL) (COMP:KDE ALL)`, the trace utility turns on all levels of tracing for all the files and functions for which KDH, ACF1, and KDE are the classes.

```
kbbcre1.c, 400, May 29 2007, 12:54:43, 1.1, *
kbbcrn1.c, 400, May 29 2007, 12:54:42, 1.1, *
kdhb1de.c, 400, May 29 2007, 12:59:34, 1.1, KDH
kdh0med.c, 400, May 29 2007, 12:59:24, 1.1, KDH
kdhsrej.c, 400, May 29 2007, 13:00:06, 1.5, KDH
kdhb1fh.c, 400, May 29 2007, 12:59:33, 1.1, KDH
kdhb1oe.c, 400, May 29 2007, 12:59:38, 1.2, KDH
kdhs1ns.c, 400, May 29 2007, 13:00:08, 1.3, KDH
kbbacd1.c, 400, May 29 2007, 12:54:27, 1.2, ACF1
kbbaclc.c, 400, May 29 2007, 12:54:27, 1.4, ACF1
kbbac1i.c, 400, May 29 2007, 12:54:28, 1.11, ACF1
vkdhscfn.c, 400, May 29 2007, 13:00:11, 1.1, KDH
kdhserq.c, 400, May 29 2007, 12:59:53, 1.1, KDH
kdhb1pr.c, 400, May 29 2007, 12:59:39, 1.1, KDH
kdhsgrh.c, 400, May 29 2007, 12:59:49, 1.1, KDH
kdh0uts.c, 400, May 29 2007, 12:59:23, 1.1, KDH
kdhsrsp.c, 400, May 29 2007, 13:00:13, 1.2, KDH
kdhs1rp.c, 400, May 29 2007, 13:00:12, 1.1, KDH
kdhscsv.c, 400, May 29 2007, 12:59:58, 1.9, KDH
kdebbac.c, 400, May 29 2007, 12:56:50, 1.10, KDE
...
```

Turning on tracing

To use the tracing utility, you must use a local logon credential for the computer. This tracing method uses the IBM Tivoli Monitoring Service Console. Access the Service Console by using a web browser.

About this task

When you start the Service Console, information is displayed about the components that are currently running on that computer. For example, these components are listed as follows:

- Tivoli Enterprise Portal Server: `cnp`
- Monitoring Agent for Windows OS: `nt`
- Tivoli Enterprise Monitoring Server: `ms`

After you log on, you can type a question mark (?) to display a list of the supported commands. Use the `ras1` command to modify trace settings. If you type this command in the field provided in the Service Console window and click **Submit**, the help for this command is displayed.

Procedure

1. Open a web browser and enter the URL to access the Service Console.

```
http://hostname:1920
```

where *hostname* is the IP address or host name of the computer on which the IBM Tivoli Monitoring component is running.

2. Click the hyperlink associated with the component for which you want to modify its trace settings.

Note: In the previous view, if you want to modify tracing for the Tivoli Enterprise Monitoring Server, select **IBM Tivoli Monitoring Service Console** under **Service Point: system.your host name_ms**.

3. Enter a user ID and password to access the system. This ID is any valid user that has access to the system.
4. Enter the command to turn on the required level of trace for the specified component classes or units.

```
ras1 set (UNIT|COMP: class_name ALL|Flow|ERROR|Detail|INPUT|Metrics|OUTPUT|STATE)
{(UNIT|COMP: class_name ALL|Flow|ERROR|Detail|INPUT|Metrics|OUTPUT|STATE)}
```

For example, to turn on the control flow trace for the KDE, the command is:

```
ras1 (COMP:KDE Flow)
```

Turning off tracing

You can use the IBM Tivoli Monitoring Service Console to run the **ras1** command and dynamically turn off tracing.

Procedure

1. Open a web browser and enter the URL to access the Service Console.

```
http://hostname:1920
```

where *hostname* is the IP address or host name of the computer on which the IBM Tivoli Monitoring component is running.

2. Click the hyperlink associated with the component for which you want to modify its trace settings.
3. Enter a user ID and password to access the system. This ID is any valid user that has access to the system.
4. Enter the command to turn off the required level of trace for the specified component classes or units.

```
ras1 set (UNIT|COMP: class_name ANY)
{(UNIT|COMP: class_name ANY)}
```

For example, to turn off tracing for the *kbbcrd* class of the Windows OS agent, the command is:

```
ras1 set (UNIT:kbbcrd ANY)
```

Setting trace parameters for the Tivoli Enterprise Console server

In addition to the trace information captured by IBM Tivoli Monitoring, you can also collect additional trace information for the Tivoli Enterprise Console components that gather event server metrics.

About this task

To collect this information, modify the `.tec_diag_config` file on the Tivoli Enterprise Console event server. Use the steps in the following procedure to modify the event server trace parameters.

Procedure

1. Open the `$BINDIR/TME/TEC/.tec_diag_config` file in an ASCII editor.
2. Locate the entries that configure trace logging for the agent components on the event server. Two entries are included, one for `tec_reception` and one for `tec_rule`:

```
# to debug Agent Utils
tec_reception Agent_Utils error /tmp/tec_reception
SP
# to debug Agent Utils
tec_rule Agent_Utils error /tmp/tec_rule
```

3. To gather additional trace information, modify these entries to specify a trace level of `trace2`:

```
# to debug Agent Utils
tec_reception Agent_Utils trace2 /tmp/tec_reception
SP
# to debug Agent Utils
tec_rule Agent_Utils trace2 /tmp/tec_rule
```

4. In addition, modify the `Highest_level` entries for `tec_rule` and `tec_reception`:

```
tec_reception Highest_level trace2
SP
tec_rule Highest_level trace2
```

Problems and workarounds

The known problems and workarounds are organized into types of problems that might occur with the Linux Kernel-based Virtual Machines agent, for example installation and configuration problems and workspace problems.

Note: You can resolve some problems by ensuring that your system matches the system requirements listed in the Prerequisites topic for the agent in the IBM Tivoli Monitoring for Virtual Environments Information Center.

For general troubleshooting information, see the *IBM Tivoli Monitoring Troubleshooting Guide*.

Installation and configuration troubleshooting

Problems can occur during installation, configuration, and uninstallation of the agent.

The problems and solutions in Table 4 can occur during installation, configuration, and uninstallation of the agent.

Table 4. Problems and solutions for installation and configuration

Problem	Solution
<p>(UNIX only) During a command-line installation, you choose to install a component that is currently installed, and you see the following warning: WARNING - you are about to install the SAME version of "component_name" where component_name is the name of the component that you are attempting to install.</p> <p>Note: This problem affects UNIX command-line installations. If you monitor only Windows environments, you see this problem if you choose to install a product component (for example, a monitoring server) on a UNIX system.</p>	<p>You must exit and restart the installation process. You cannot return to the list where you selected components to install. When you run the installer again, do not attempt to install any component that is currently installed.</p>
<p>A message similar to "Unable to find running CMS on CT_CMSLIST" in the log file is displayed.</p>	<p>If a message similar to "Unable to find running CMS on CT_CMSLIST" is displayed in the log file, the agent cannot connect to the monitoring server. Confirm the following points:</p> <ul style="list-style-type: none">• Do multiple network interface cards (NICs) exist on the system?• If multiple NICs exist on the system, find out which one is configured for the monitoring server. Ensure that you specify the correct host name and port settings for communication in the IBM Tivoli Monitoring environment.
<p>The system is experiencing high CPU usage.</p>	<p>Agent process: View the memory usage of the KV1CMA process. If CPU usage seems to be excessive, restart the monitoring agent.</p> <p>Network cards: The network card configurations can decrease the performance of a system. Each stream of packets that a network card receives (assuming that it is a broadcast or destined for the under-performing system) must generate a CPU interrupt and transfer the data through the I/O bus. If the network card in question is a bus-mastering card, work can be offloaded and a data transfer between memory and the network card can continue without using CPU processing power. Bus-mastering cards are 32-bit and are based on PCI or EISA bus architectures.</p>

Table 4. Problems and solutions for installation and configuration (continued)

Problem	Solution
<p>The configuration panel is blank on 64-bit Windows systems where the Tivoli Enterprise Monitoring Agent Framework (component GL) is version 06.23.00.00 or 06.23.01.00.</p>	<p>Check the GL component version by running <code>kincinfo -t GL</code> from a Windows command line. Example: <code>%CANDLE_HOME%\Install\ITM\kincinfo -t GL</code></p> <p>If the GL component version is 06.23.00.00 or 06.23.01.00, take one of the following actions:</p> <ul style="list-style-type: none"> • Preferred action: Upgrade the Windows OS Agent to Version 6.2.3 Fix Pack 2. • Alternate action: Install the Agent Compatibility (AC) component from the IBM Tivoli Monitoring V6.2.3 Fix Pack 1 media. See Installing the Agent Compatibility (AC) component (http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.3fp1/itm623FP1_install199.htm?path=3_0_3_0_6_3_3_0_0#acpinstall).

Table 5. General problems and solutions for uninstallation

Problem	Solution
<p>The way to remove inactive managed systems (systems whose status is OFFLINE) from the Navigator tree in the portal is not obvious.</p>	<p>Use the following steps to remove, but not uninstall, an offline managed system from the Navigator tree:</p> <ol style="list-style-type: none"> 1. Click the Enterprise icon in the Navigator tree. 2. Right-click, and then click Workspace > Managed System Status. 3. Right-click the offline managed system, and select Clear offline entry. <p>To uninstall the monitoring agent, use the procedure described in the <i>IBM Tivoli Monitoring Installation and Setup Guide</i>.</p>

Table 5. General problems and solutions for uninstallation (continued)

Problem	Solution
<p>IBM Tivoli Monitoring might not be able to generate a unique name for monitoring components because of the truncation of names that the product automatically generates.</p>	<p>If the agent supports multiple instances, IBM Tivoli Monitoring automatically creates a name for each monitoring component by concatenating the subsystem name, host name, and product code separated by colons (<i>subsystem_name:hostname:KV1</i>).</p> <p>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.</p> <p>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 <i>hostname</i> portion of the name as follows:</p> <ol style="list-style-type: none"> 1. Open the configuration file for the monitoring agent, which is located in the following path: <ul style="list-style-type: none"> • On UNIX and Linux: itm_home/config/<i>product_code.ini</i> and <i>product_code.config</i>. For example, the file names for the Monitoring Agent for UNIX OS is <i>ux.ini</i> and <i>ux.config</i>. 2. Find the line that begins with CTIRA_HOSTNAME=. 3. Type a new name for host name that is a unique, shorter name for the host computer. The final concatenated name including the subsystem name, new host name, and KV1, cannot be longer than 32 characters. <p>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.</p> 4. Save the file. 5. Restart the agent.
<p>The software inventory tag for the agent on UNIX and Linux systems is not removed during uninstallation of the agent.</p>	<p>After uninstalling the agent, manually remove the file named <i>full name of agent.cmptag</i> from the \$CANDLEHOME/properties/version/ directory.</p>

Table 5. General problems and solutions for uninstallation (continued)

Problem	Solution
<p>Installation on RHEL Linux 64-bit systems uses install.sh. Running this script fails with a runGSkit failure: Return error code: 99.</p>	<p>GSkit is called by install.sh and fails when runGSkit calls verifyInstall. Review the <i>InstallDirectory/logs/candle_installation.log</i> file and look for references to runGSkit.</p> <p>For example, output similar to the following might be present:</p> <pre>runGSkit: ----- Running command: /opt/IBM/ITM/li6243/gs/bin/private_verifyinstall /opt/IBM/ITM/li6243/gs/bin/gsk7ver: error while loading shared libraries: libstdc++.so.5: cannot open shared object file: No such file or directory Error: Verify Failed Expected Details of gskit in /opt/IBM/ITM/li6243/gs runGSkit: return code from command is 99 runGSkit: ----- End of running command ----- runGSkit: error Return error code: 99 runGSkit: error GSKit check failure, script: /opt/IBM/ITM/li6243/gs/bin/ private_verifyinstall runGSkit: error li6243 - GSK check error, verifyInstall test failed</pre> <p>In the previous example, the 32-bit version of the libstdc++.so.5 file is not present. This file comes from the compat-libstdc++-33-3.2.3-XX.i686.rpm package, which is not installed on 64-bit RHEL systems by default. When this package is installed, the problem no longer occurs.</p>
<p>After installation, the Linux Kernel-based Virtual Machines agent instance fails to start. The following message appears in the agent log:</p> <pre>(4CF55620.003F-1:kbbssge.c,52,"BSS1_GetEnv") KBB_SIG1="-asyncoff -syncoff -dumpoff" (4CF55620.0040-1:signalmanager.cpp,170, "startManagerThread") Error starting signal manager thread. Return code = 11; Resource temporarily unavailable. Use the return code and message to investigate the failure. Agent is terminating.</pre>	<p>The probable cause of the problem is the public domain Korn shell, pdksh. Uninstall the pdksh shell and install the ksh rpm that is included on the Linux installation media.</p>

Remote deployment troubleshooting

Problems can occur with remote deployment and removal of agent software using the Agent Remote Deploy process.

Table 6 contains problems and solutions related to remote deployment.

Table 6. Remote deployment problems and solutions

Problem	Solution
<p>While you are using the remote deployment feature to install the IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines, an empty command window is displayed on the target computer. This problem occurs when the target of remote deployment is a Windows computer. (For more information about the remote deployment feature, see the <i>IBM Tivoli Monitoring Installation and Setup Guide</i>.)</p>	<p>Do not close or modify this window. It is part of the installation process and is dismissed automatically.</p>

Table 6. Remote deployment problems and solutions (continued)

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 occur when you attempt the remote removal process immediately after you restart 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.

Agent troubleshooting

A problem can occur with the agent after it has been installed.

Table 7 contains problems and solutions that can occur with the agent after it is installed.

Table 7. Agent problems and solutions

Problem	Solution
Log data accumulates too rapidly.	Check the RAS trace option settings, which are described in "Setting RAS trace parameters by using the GUI" on page 71. The trace option settings that you can set on the KBB_RAS1= and KDC_DEBUG= lines potentially generate large amounts of data.
When using the itmcmd agent commands to start or stop this monitoring agent, you receive the following error message: MKCIIN0201E Specified product is not configured.	Include the command option -o to specify the instance to start or stop. The instance name must match the name used for configuring the agent. For example: ./itmcmd agent -o Test1 start v1 For more information about using the itmcmd commands, see the <i>IBM Tivoli Monitoring Command Reference</i> .

Table 7. Agent problems and solutions (continued)

Problem	Solution
<p>A configured and running instance of the monitoring agent is not displayed in the Tivoli Enterprise Portal, but other instances of the monitoring agent on the same system are displayed in the portal.</p>	<p>IBM Tivoli Monitoring products use Remote Procedure Call (RPC) to define and control product behavior. RPC is the mechanism that a client process uses 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 protocol (or delivery mechanism) for RPCs that you want.</p> <p>IP.PIPE is the name given to Tivoli TCP/IP protocol for RPCs. The RPCs are socket-based operations that use TCP/IP ports to form socket addresses. IP.PIPE implements virtual sockets and multiplexes all virtual socket traffic across a single physical TCP/IP port (visible from the netstat command).</p> <p>A Tivoli process derives the physical port for IP.PIPE communications based on the configured, well-known port for the hub Tivoli Enterprise Monitoring Server. (This well-known port or BASE_PORT is configured by using the 'PORT:' keyword on the KDC_FAMILIES / KDE_TRANSPORT environment variable and defaults to '1918'.)</p> <p>The physical port allocation method is defined as (BASE_PORT + 4096*N), where N=0 for a Tivoli Enterprise Monitoring Server process and N={1, 2, ..., 15} for another type of monitoring server process. Two architectural limits result as a consequence of the physical port allocation method:</p> <ul style="list-style-type: none"> • No more than one Tivoli Enterprise Monitoring Server reporting to a specific Tivoli Enterprise Monitoring Server hub can be active on a system image. • No more than 15 IP.PIPE processes can be active on a single system image. <p>A single system image can support any number of Tivoli Enterprise Monitoring Server processes (address spaces) if each Tivoli Enterprise Monitoring Server on that image reports to a different hub. By definition, one Tivoli Enterprise Monitoring Server hub is available per monitoring enterprise, so this architecture limit has been reduced to one Tivoli Enterprise Monitoring Server per system image.</p> <p>No more than 15 IP.PIPE processes or address spaces can be active on a single system image. With the first limit expressed earlier, this second limitation refers specifically to Tivoli Enterprise Monitoring Agent processes: no more than 15 agents per system image.</p> <p>Continued on next row.</p>

Table 7. Agent problems and solutions (continued)

Problem	Solution
Continued from previous row.	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 the EPHEMERAL IP.PIPE process. (This process is IP.PIPE configured with the 'EPHEMERAL:Y' keyword in the KDC_FAMILIES / KDE_TRANSPORT environment variable). The number of ephemeral IP.PIPE connections per system image has no limitation. If ephemeral endpoints are used, the Warehouse Proxy agent is accessible from the Tivoli Enterprise Monitoring Server associated with the agents using ephemeral connections either by running the Warehouse Proxy agent on the same computer or by using the Firewall Gateway feature. (The Firewall Gateway feature relays the Warehouse Proxy agent connection from the Tivoli Enterprise Monitoring Server computer to the Warehouse Proxy agent computer if the Warehouse Proxy agent cannot coexist on the same computer.)
In the agent log, you see the message SEVERE: DataSource.connect: failed to connect to data source <i>ip address</i> .	Find the hypervisor URI of the host that failed to connect. It is listed in a message just before the failed to connect message. On the computer where the agent is installed, enter the command <code>export LIBVIRT_DEBUG=yes</code> followed by <code>virsh -chypervisor uri that failed</code> . Review the extra debug messages for symptoms of an underlying problem.
After turning on extra libvirt messages by using <code>export LIBVIRT_DEBUG=yes</code> , when you execute the command <code>virsh -chypervisor uri</code> , you see the message: <code>bash: nc: command not found</code> .	Install netcat-openbsd on the host of the hypervisor to be monitored.

Workspace troubleshooting

Problems can occur with general workspaces and agent-specific workspaces.

Table 8 on page 84 contains problems and solutions related to workspaces.

Table 8. Workspace problems and solutions

Problem	Solution
<p>The process application components are available, but the Availability status shows PROCESS_DATA_NOT_AVAILABLE.</p>	<p>This problem occurs because the PerfProc performance object is disabled. When this condition exists, IBM Tivoli Monitoring cannot collect performance data for this process. Use the following steps to confirm that this problem exists and to resolve it:</p> <ol style="list-style-type: none"> 1. In the Windows Start menu, click Run. 2. Type perfmon.exe in the Open field of the Run window. The Performance window is displayed. 3. Click the plus sign (+) in the toolbar. The Add Counters window is displayed. 4. Look for Process in the Performance object menu. 5. Complete one of the following actions: <ul style="list-style-type: none"> • If you see Process in the menu, the PerfProc performance object is enabled and the problem is coming from a different source. You might need to contact IBM Software Support. • If you do not see Process in the menu, use the Microsoft utility from the Microsoft.com Operations website to enable the PerfProc performance object. The Process performance object becomes visible in the Performance object menu of the Add Counters windows, and IBM Tivoli Monitoring is able to detect Availability data. 6. Restart the monitoring agent.
<p>The name of the attribute does not display in a bar chart or graph view.</p>	<p>When a chart or graph view that includes the attribute is scaled to a small size, a blank space is displayed instead of a truncated name. To see the name of the attribute, expand the view of the chart until sufficient space is available to display all characters of the attribute name.</p>
<p>At the end of each view, you see the following Historical workspace KFWITM220E error: Request failed during execution.</p>	<p>Ensure that you configure all groups that supply data to the view. In the Historical Configuration view, ensure that data collection is started for all groups that supply data to the view.</p>

Table 8. Workspace problems and solutions (continued)

Problem	Solution
<p>You start collection of historical data but the data cannot be seen.</p>	<p>Use the following managing options for historical data collection:</p> <ul style="list-style-type: none"> • Basic historical data collection populates the Warehouse with raw data. This type of data collection is turned off by default. For information about managing this feature including how to set the interval at which data is collected, see <i>Managing historical data</i> in the <i>IBM Tivoli Monitoring Administrator's Guide</i>. By setting a more frequent interval for data collection, you reduce the load on the system incurred every time data is uploaded. • Use the Summarization and Pruning agent to collect specific amounts and types of historical data. Historical data is not displayed until the Summarization and Pruning monitoring agent begins collecting the data. By default, this agent begins collection at 2 a.m. daily. At that point, data is visible in the workspace view. For information about how to modify the default collection settings, see <i>Managing historical data</i> in the <i>IBM Tivoli Monitoring Administrator's Guide</i>.
<p>Historical data collection is unavailable because of incorrect queries in the Tivoli Enterprise Portal.</p>	<p>The Sort By, Group By, and First/Last functions column are not compatible with the historical data collection feature. Use of these advanced functions makes a query ineligible for historical data collection.</p> <p>Even if data collection has started, you cannot use the time span feature if the query for the chart or table includes column functions or advanced query options (Sort By, Group By, First / Last).</p> <p>To ensure support of historical data collection, do not use the Sort By, Group By, or First/Last functions in your queries.</p> <p>For information about the historical data collection function, See <i>Managing historical data</i> in the <i>IBM Tivoli Monitoring Administrator's Guide</i> or the Tivoli Enterprise Portal online help .</p>
<p>When you use a long process name in the situation, the process name is truncated.</p>	<p>Truncation of process or service names for situations in the Availability table in the portal display is the expected behavior. The maximum name length is 100 bytes.</p>
<p>Regular (non-historical) monitoring data fails to be displayed.</p>	<p>Check the formation of the queries you use to gather data. For example, look for invalid SQL statements.</p>
<p>Navigator items and workspace titles are labeled with internal names such as Kxx:KXX0000 instead of the correct names (such as Disk), where XX and xx represent the two-character agent code.</p>	<p>Ensure that application support has been added on the monitoring server, portal server, and portal client.</p> <p>For more information about installing application support, see <i>Installing and enabling application support</i> in the <i>IBM Tivoli Monitoring Installation and Setup Guide</i>.</p>

Situation troubleshooting

Problems can occur with situations and situation configuration.

Table 9 contains problems and solutions for situations.

Table 9. Situation problems and solutions

Problem	Solution
Monitoring activity requires too much disk space.	Check the RAS trace logging settings that are described in “Setting RAS trace parameters by using the GUI” on page 71. For example, trace logs grow rapidly when you apply the ALL logging option.
Monitoring activity requires too many system resources.	“Disk capacity planning for historical data” on page 51 describes the performance impact of specific attribute groups. If possible, decrease your use of the attribute groups that require greater system resources.
A formula that uses mathematical operators appears to be incorrect. For example, if you were monitoring a Linux system, the formula that calculates when Free Memory falls under 10 percent of Total Memory does not work: <code>LT #'Linux_VM_Stats.Total_Memory' / 10</code>	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. In the example, you can select the % Memory Free attribute and avoid the need for math operators.
You want to change the appearance of situations when they are displayed in the navigation tree.	<ol style="list-style-type: none"> 1. Right-click an item in the navigation tree. 2. Click Situations in the menu. The Situation Editor window is displayed. 3. Select the situation that you want to modify. 4. Use the State menu to set the status and appearance of the Situation when it triggers. Note: The State setting is not related to severity settings in the Tivoli Enterprise Console.
When a situation is triggered in the Event Log attribute group, it remains in the Situation Event Console as long as the event ID entry is present in the Event Log workspace. When this event ID entry is removed from the Event Log workspace on the Tivoli Enterprise Portal, the situation is also cleared even if the actual problem that caused the event is not resolved, and the event ID entry is also present in the Windows Event Viewer.	<p>A timeout occurs on the cache of events for the NT Event Log group. Increase the cache time of Event Log collection to meet your requirements by adding the following variable and timeout value to the <code>KpcENV</code> file for the agent (where <code>pc</code> is the two-letter product code): <code>CDP_NT_EVENT_LOG_CACHE_TIMEOUT=3600</code></p> <p>This variable determines how long events from the NT Event Log are kept.</p>
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 not displayed, confirm that the monitoring server has been seeded for the agent. If not, seed the server, as described in the <i>IBM Tivoli Monitoring Installation and Setup Guide</i> .
The monitoring interval is too long.	Access the Situation Editor view for the situation that you want to modify. Check the Sampling interval area in the Formula tab. Adjust the time interval as required.
The situation did not activate at startup.	<p>Manually recycle the situation as follows:</p> <ol style="list-style-type: none"> 1. Right-click the situation and select Stop Situation. 2. Right-click the situation and select Start Situation. Note: You can permanently avoid this problem by selecting the Run at Startup check box of the Situation Editor view for a specific situation.

Table 9. Situation problems and solutions (continued)

Problem	Solution
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 did not occur even though the predicate was correctly specified.	Check the logs, reports, and workspaces.
A situation fires on an unexpected managed object.	Confirm that you distributed and started the situation on the correct managed system.
The product did not distribute the situation to a managed system.	Click the Distribution tab and check the distribution settings for the situation.
The situation does not fire.	<p>This problem can be caused when 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.</p> <p>In the Formula tab, analyze predicates as follows:</p> <ol style="list-style-type: none"> 1. Click the fx icon in the Formula area. The Show formula window is displayed. <ol style="list-style-type: none"> a. Confirm the following details in the Formula area of the window: <ul style="list-style-type: none"> • The attributes that you intend to monitor are specified in the formula. • The situations that you intend to monitor are specified in the formula. • The logical operators in the formula match your monitoring goal. • The numeric values in the formula match your monitoring goal. b. (Optional) Select the Show detailed formula check box 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 numeric values that immediately trigger a monitoring event. The triggering of the event confirms that other predicates in the formula are valid. <p>Note: After you complete this test, you must restore the numeric values to valid levels so that you do not generate excessive monitoring data based on your temporary settings.</p> <p>For additional information about situations that do not fire, see <i>Situations are not firing</i> in the <i>IBM Tivoli Monitoring Troubleshooting Guide</i>.</p>
Situation events are not displayed in the Events Console view of the workspace.	Associate the situation with a Navigator item. Note: The situation does not need to be displayed in the workspace. It is sufficient that the situation is associated with any Navigator item.

Table 9. Situation problems and solutions (continued)

Problem	Solution
You do not have access to a situation.	<p>Note: You must have administrator privileges to complete these steps.</p> <ol style="list-style-type: none"> 1. Click 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 role. 4. Click OK.
A managed system seems to be offline.	<ol style="list-style-type: none"> 1. Select Physical View and click the Enterprise Level of the navigator tree. 2. Click 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 the status of the specific system or application.

Take Action commands troubleshooting

Problems can occur with Take Action commands.

Table 10 contains problems and solutions that can occur with Take Action commands.

When each Take Action command runs, it generates a log file listed in Table 3 on page 67.

Table 10. Take Action commands problems and solutions

Problem	Solution
Take Action commands often require several minutes to complete.	Allow several minutes. If you do not see a message advising you of completion, try to run the command manually.
Situations fail to trigger Take Action commands.	Attempt to manually run the Take Action command in the Tivoli Enterprise Portal. If the Take Action command works, look for configuration problems in the situation. See "Situation troubleshooting" on page 86. If the Take Action command fails, for general information about troubleshooting Take Action commands, see the <i>IBM Tivoli Monitoring Troubleshooting Guide</i> .

Support information

If you have a problem with your IBM software, you want to resolve it quickly.

IBM provides the following ways for you to obtain the support you need:

Online

The following websites contain troubleshooting information:

- Go to the IBM Software Support website (<http://www.ibm.com/support/entry/portal/software>) and follow the instructions.
- Go to the Application Performance Management Wiki (<http://www.ibm.com/developerworks/servicemanagement/apm/index.html>). Feel free to contribute to this wiki.

IBM Support Assistant

The IBM Support Assistant (ISA) is a free local software serviceability workbench that helps you resolve questions and problems with IBM software products. The ISA provides quick access to support-related information and serviceability tools for problem determination. To install the ISA software, go to the IBM Support Assistant website (<http://www.ibm.com/software/support/isa>).

Informational, warning, and error messages overview

Messages relay information about how the system or application is performing and can alert you to exceptional conditions when they occur.

Messages are sent to an output destination, such as a file, database, or console screen.

If you receive a warning or error message, you can do one of the following actions:

- Follow the instructions listed in the Detail window of the message if this information is included there.
- Consult the message details listed in this topic to see what action you can take to correct the problem.
- Consult the message log for message ID, text, time, and date of the message, as well as other data you can use to diagnose the problem.

Message format

The message format contains a message ID and text, an explanation, and an operator response.

IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines messages have the following format:

Message ID and text
Explanation
Operator Response

The message ID has the following format:

`CCC###severity`

where:

CCC Prefix that indicates the component to which the message applies. The following components are used:

KV1 General Linux Kernel-based Virtual Machines agent messages

Number of the message

severity

Severity of the message. Three levels of severity are used:

- I** Informational messages provide feedback about something that happened in the product or system that might be important. These messages can provide guidance when you are requesting a specific action from the product.
- W** Warning messages call your attention to an exception condition. The condition might not be an error but can cause problems if not resolved.
- E** Error messages indicate that an action cannot be completed because of a user or system error. These messages require user response.

The *Text* of the message provides a general statement regarding the problem or condition that occurred. The *Explanation* provides additional information about the message and the possible cause for the condition. The *Operator Response* provides actions to take in response to the condition, particularly for error messages (messages with the "E" suffix).

Note: Many message texts and explanations contain variables, such as the specific name of a server or application. Those variables are represented in this topic as symbols, such as "&1." Actual messages contain values for these variables.

Agent messages

The following messages apply to IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines.

Appendix A. Event mapping

The Tivoli Event Integration Facility (EIF) interface is used to forward situation events to Tivoli Netcool/OMNIBus or Tivoli Enterprise Console.

EIF events specify an event class, and the event data is specified as name-value pairs that identify the name of an event slot and the value for the slot. An event class can have subclasses. IBM Tivoli Monitoring provides the base event class definitions and a set of base slots that are included in all monitoring events. Agents extend the base event classes to define subclasses that include agent-specific slots. For Linux Kernel-based Virtual Machines agent events, the event classes correspond to the agent attribute groups, and the agent-specific slots correspond to the attributes in the attribute group.

The situation editor in the Tivoli Enterprise Portal can be used to perform custom mapping of data to EIF slots instead of using the default mapping described in this topic. For more information about EIF slot customization, see the *Tivoli Enterprise Portal User's Guide*.

Tivoli Enterprise Console requires that event classes and their slots are defined in BAROC (Basic Recorder of Objects in C) files. Each agent provides a BAROC file that contains event class definitions for the agent and is installed on the Tivoli Enterprise Monitoring Server in the TECLIB directory (`install_dir/cms/TECLIB` for Windows systems and `install_dir/tables/TEMS_hostname/TECLIB` for UNIX systems) when application support for the agent is installed. The BAROC file for the agent and the base BAROC files provided with Tivoli Monitoring must also be installed onto the Tivoli Enterprise Console. For details, see "Setting up event forwarding to Tivoli Enterprise Console" in the *IBM Tivoli Monitoring Installation and Setup Guide*.

Each of the event classes is a child of KV1_Base and is defined in the `kv1.baroc` (version 07.20.00) file. The KV1_Base event class can be used for generic rules processing for any event from the IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines.

For events that are generated by situations in the Host CPU attribute group, events are sent by using the ITM_KV1_HOST_CPU event class. This event class contains the following slots:

- node: STRING
- timestamp: STRING
- host_name: STRING
- host_name_enum: STRING
- number_of_cpus: REAL
- number_of_cpus_enum: STRING
- number_of_cpu_ghz: REAL
- number_of_cpu_ghz_enum: STRING
- vm_cpus_consumed: REAL
- vm_cpus_consumed_enum: STRING
- vm_cpu_ghz_consumed: REAL
- vm_cpu_ghz_consumed_enum: STRING
- cpus_unaccounted_for: REAL
- cpus_unaccounted_for_enum: STRING
- cpu_ghz_unaccounted_for: REAL
- cpu_ghz_unaccounted_for_enum: STRING
- cpus_consumed_per_vm: REAL

- cpus_consumed_per_vm_enum: STRING
- cpu_ghz_consumed_per_vm: REAL
- cpu_ghz_consumed_per_vm_enum: STRING
- number_of_active_vms: REAL
- number_of_active_vms_enum: STRING
- cpu_percent_consumed: REAL
- cpu_percent_consumed_enum: STRING
- cpu_percent_unaccounted: REAL
- cpu_percent_unaccounted_enum: STRING
- vm_cpu_fit_estimate: REAL
- vm_cpu_fit_estimate_enum: STRING
- virtual_cpus_allocated: REAL
- virtual_cpus_allocated_enum: STRING
- over_under_allocated_cpu: REAL
- over_under_allocated_cpu_enum: STRING
- over_under_allocated_cpu_ghz: REAL
- over_under_allocated_cpu_ghz_enum: STRING
- percent_cpu_allocated: REAL
- percent_cpu_allocated_enum: STRING

For events that are generated by situations in the Host Memory attribute group, events are sent by using the ITM_KV1_HOST_MEMORY event class. This event class contains the following slots:

- node: STRING
- timestamp: STRING
- host_name: STRING
- host_name_enum: STRING
- memory_size: REAL
- memory_size_enum: STRING
- vm_memory_allocated: REAL
- vm_memory_allocated_enum: STRING
- vm_memory_used: REAL
- vm_memory_used_enum: STRING
- over_under_allocated: REAL
- over_under_allocated_enum: STRING
- over_under_used: REAL
- over_under_used_enum: STRING
- percent_memory_allocated: REAL
- percent_memory_allocated_enum: STRING
- percent_memory_used: REAL
- percent_memory_used_enum: STRING
- percent_memory_unallocated: REAL
- percent_memory_unallocated_enum: STRING
- percent_memory_unused: REAL
- percent_memory_unused_enum: STRING
- number_of_active_vms: REAL

- number_of_active_vms_enum: STRING
- memory_allocated_per_vm: REAL
- memory_allocated_per_vm_enum: STRING
- memory_used_per_vm: REAL
- memory_used_per_vm_enum: STRING
- vm_memory_allocated_fit_estimate: REAL
- vm_memory_allocated_fit_estimate_enum: STRING
- vm_memory_used_fit_estimate: REAL
- vm_memory_used_fit_estimate_enum: STRING

For events that are generated by situations in the Hosts attribute group, events are sent by using the ITM_KV1_HOSTS event class. This event class contains the following slots:

- node: STRING
- timestamp: STRING
- host_name: STRING
- host_name_enum: STRING
- ip_address: STRING
- ip_address_enum: STRING
- hypervisor_uri: STRING
- hypervisor_uri_enum: STRING
- protocol: STRING
- protocol_enum: STRING
- cpu_model: STRING
- cpu_model_enum: STRING
- number_of_cpus: REAL
- number_of_cpus_enum: STRING
- cpu_frequency: REAL
- cpu_frequency_enum: STRING
- number_of_cpu_ghz: REAL
- number_of_cpu_ghz_enum: STRING
- number_of_nodes: REAL
- number_of_nodes_enum: STRING
- sockets_per_node: REAL
- sockets_per_node_enum: STRING
- cores_per_socket: REAL
- cores_per_socket_enum: STRING
- threads_per_core: REAL
- threads_per_core_enum: STRING
- memory_size: REAL
- memory_size_enum: STRING
- number_of_vms: REAL
- number_of_vms_enum: STRING
- number_of_active_vms: REAL
- number_of_active_vms_enum: STRING

For events that are generated by situations in the Performance Object Status attribute group, events are sent by using the ITM_KV1_PERFORMANCE_OBJECT_STATUS event class. This event class contains the following slots:

- node: STRING
- timestamp: STRING
- query_name: STRING
- object_name: STRING
- object_type: INTEGER
- object_type_enum: STRING
- object_status: INTEGER
- object_status_enum: STRING
- error_code: INTEGER
- error_code_enum: STRING
- last_collection_start: STRING
- last_collection_start_enum: STRING
- last_collection_finished: STRING
- last_collection_finished_enum: STRING
- last_collection_duration: REAL
- average_collection_duration: REAL
- average_collection_duration_enum: STRING
- refresh_interval: INTEGER
- number_of_collections: INTEGER
- cache_hits: INTEGER
- cache_misses: INTEGER
- cache_hit_percent: REAL
- intervals_skipped: INTEGER

For events that are generated by situations in the Scheduler Parameters attribute group, events are sent by using the ITM_KV1_SCHEDULER_PARAMETERS event class. This event class contains the following slots:

- node: STRING
- timestamp: STRING
- host_name: STRING
- host_name_enum: STRING
- virtual_machine_name: STRING
- virtual_machine_name_enum: STRING
- scheduler_parameter_name: STRING
- scheduler_parameter_name_enum: STRING
- scheduler_parameter_value: STRING
- scheduler_parameter_value_enum: STRING
- scheduler_parameter_type: STRING
- scheduler_parameter_type_enum: STRING

For events that are generated by situations in the Storage Pools attribute group, events are sent by using the ITM_KV1_STORAGE_POOLS event class. This event class contains the following slots:

- node: STRING
- timestamp: STRING

- host_name: STRING
- host_name_enum: STRING
- storage_pool_name: STRING
- storage_pool_name_enum: STRING
- storage_pool_uuid: STRING
- storage_pool_uuid_enum: STRING
- storage_pool_state: STRING
- storage_pool_state_enum: STRING
- storage_pool_type: STRING
- storage_pool_type_enum: STRING
- storage_pool_capacity: REAL
- storage_pool_capacity_enum: STRING
- storage_pool_used: REAL
- storage_pool_used_enum: STRING
- storage_pool_available: REAL
- storage_pool_available_enum: STRING
- percent_used: REAL
- percent_used_enum: STRING

For events that are generated by situations in the Virtual Machines attribute group, events are sent by using the ITM_KV1_VIRTUAL_MACHINES event class. This event class contains the following slots:

- node: STRING
- timestamp: STRING
- host_name: STRING
- host_name_enum: STRING
- virtual_machine_name: STRING
- virtual_machine_name_enum: STRING
- virtual_machine_uuid: STRING
- virtual_machine_uuid_enum: STRING
- virtual_machine_state: STRING
- virtual_machine_state_enum: STRING
- virtualization_type: STRING
- virtualization_type_enum: STRING
- number_of_virtual_cpus: REAL
- number_of_virtual_cpus_enum: STRING
- sample_timestamp: STRING
- sample_timestamp_enum: STRING
- sample_time_delta: REAL
- sample_time_delta_enum: STRING
- cpu_time: REAL
- cpu_time_enum: STRING
- cpu_time_delta: REAL
- cpu_time_delta_enum: STRING
- cpus_consumed: REAL
- cpus_consumed_enum: STRING

- cpu_percent: REAL
- cpu_percent_enum: STRING
- memory_allocated: REAL
- memory_allocated_enum: STRING
- memory_used: REAL
- memory_used_enum: STRING
- memory_percent: REAL
- memory_percent_enum: STRING
- action_on_poweroff: STRING
- action_on_poweroff_enum: STRING
- action_on_reboot: STRING
- action_on_reboot_enum: STRING
- action_on_crash: STRING
- action_on_crash_enum: STRING
- cpu_match: STRING
- cpu_match_enum: STRING
- cpu_model: STRING
- cpu_model_enum: STRING
- number_of_sockets: REAL
- number_of_sockets_enum: STRING
- cores_per_socket: REAL
- cores_per_socket_enum: STRING
- threads_per_core: REAL
- threads_per_core_enum: STRING
- cpu_shares: REAL
- cpu_shares_enum: STRING

Appendix B. Documentation library

Various publications are relevant to the use of the IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines.

For information about how to access and use the publications, see *Using the publications* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/common/using_publications.htm).

To find publications from the previous version of a product, click **Previous versions** under the name of the product in the **Contents** pane.

IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines library

The documentation for this agent and other product components is in the IBM Tivoli Monitoring for Virtual Environments Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.tivoli.itmvs.doc_7.2/welcome_ve72.htm).

One document is specific to the Linux Kernel-based Virtual Machines agent. The IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines User's Guide provides agent-specific information for configuring, using, and troubleshooting the Linux Kernel-based Virtual Machines agent.

The **Prerequisites** topic in the information center contains information about the prerequisites for each component.

Use the information in the user's guide for the agent with the *Tivoli Enterprise Portal User's Guide* to monitor Linux Kernel-based Virtual Machines resources.

Prerequisite publications

To use the information in this publication effectively, you must have some prerequisite knowledge.

See the following publications to gain the required prerequisite knowledge:

- *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 High Availability Guide for Distributed Systems*
- *IBM Tivoli Monitoring: Messages*
- *IBM Tivoli Monitoring Troubleshooting Guide*
- *IBMTivoli Monitoring: IBM i OS Agent User's Guide*
- *IBM Tivoli Monitoring: Linux OS Agent User's Guide*
- *IBM Tivoli Monitoring: UNIX OS Agent User's Guide*
- *IBM Tivoli Monitoring: Windows OS Agent User's Guide*
- *Tivoli Enterprise Portal User's Guide*
- *IBM Tivoli Performance Analyzer User's Guide*
- *IBM Tivoli Warehouse Proxy Agent User's Guide*

- *IBM Tivoli Warehouse Summarization and Pruning Agent User's Guide*
-

Related publications

The publications in related information centers provide useful information.

See the following information centers, which you can find by accessing Tivoli Documentation Central (<http://www.ibm.com/tivoli/documentation>):

- Tivoli Monitoring
 - Tivoli Application Dependency Discovery Manager
 - Tivoli Business Service Manager
 - Tivoli Common Reporting
 - Tivoli Enterprise Console
-

Other sources of documentation

You can obtain additional technical documentation about monitoring products from other sources.

See the following sources of technical documentation about monitoring products:

- Service Management Connect (SMC)

For introductory information about SMC, see IBM Service Management Connect (<http://www.ibm.com/developerworks/servicemanagement/>).

For information about Tivoli products, see the Application Performance Management community on SMC (<http://www.ibm.com/developerworks/servicemanagement/apm/index.html>).

Connect, learn, and share with Service Management professionals. Get access to developers and product support technical experts who provide their perspectives and expertise. You can use SMC for these purposes:

- Become involved with transparent development, an ongoing, open engagement between external users and developers of Tivoli products where you can access early designs, sprint demos, product roadmaps, and pre-release code.
- Connect one-on-one with the experts to collaborate and network about Tivoli and Integrated Service Management.
- Benefit from the expertise and experience of others using blogs.
- Collaborate with the broader user community using wikis and forums.
- IBM Integrated Service Management Library (<http://www.ibm.com/software/brandcatalog/ismlibrary/>) is an online catalog that contains integration documentation as well as other downloadable product extensions.
- IBM Redbook publications (<http://www.redbooks.ibm.com/>) include Redbooks® publications, Redpapers, and Redbooks technotes that provide information about products from platform and solution perspectives.
- Technotes (<http://www.ibm.com/support/entry/portal/software>), which are found through the IBM Software Support website, provide the latest information about known product limitations and workarounds.

Notices

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

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

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

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

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

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

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

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

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

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

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

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

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

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

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

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

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

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

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

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

© IBM 2009. Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2009. All rights reserved.

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

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and trademark information (www.ibm.com/legal/copytrade.shtml).

Intel, Intel logo, and Intel Xeon, 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 Oracle and/or its affiliates.

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

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

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

Index

A

- Action On Crash attribute 49
- Action On PowerOff attribute 48
- Action On Reboot attribute 49
- activities 63
- additional information
 - attributes 21
 - situations 53
 - Take Action commands 61
 - Workspaces 15
- agent
 - functions 1
 - problems and workarounds 81
- Agent Management Services 3
- attribute group 22
- attribute groups
 - Host CPU 22
 - Host Memory 27
 - Hosts 31
 - list of all 21
 - overview 21
 - Performance Object Status 35
 - Scheduler Parameters 40
 - Storage Pools 42
 - Virtual Machines 44
- attributes 22
 - Action On Crash 49
 - Action On PowerOff 48
 - Action On Reboot 49
 - additional information 21
 - Average Collection Duration 39
 - Cache Hit Percent 40
 - Cache Hits 39
 - Cache Misses 39
 - Cores Per Socket 34, 50
 - CPU Frequency (GHz) 33
 - CPU GHz Consumed Per VM 25
 - CPU GHz Unaccounted For 24
 - CPU Match 49
 - CPU Model 32, 49
 - CPU Percent 47
 - CPU Percent Consumed by VMs 25
 - CPU Percent Unaccounted 25
 - CPU Shares 50
 - CPU Time Delta (sec) 47
 - CPUs Consumed 47
 - CPUs Consumed Per VM 24
 - CPUs Unaccounted For 24
 - Error Code 37
 - Host CPU 22
 - Host Memory 27
 - Host Name 23, 27, 31, 40, 42, 45
 - Hosts 31
 - Hypervisor URI 32
 - Intervals Skipped 40
 - IP Address 32
 - Last Collection Duration 39
 - Last Collection Finished 38
 - Last Collection Start 38
 - Memory Allocated (GB) 48
 - Memory Allocated Per VM 30

- attributes (*continued*)
 - Memory Percent 48
 - Memory Size 28
 - Memory Size (GB) 34
 - Memory Used (GB) 48
 - Memory Used Per VM 30
 - Node 22, 27, 31, 35, 40, 42, 44
 - Number of Active VMs 25, 30, 35
 - Number of Collections 39
 - Number of CPU GHz 23, 33
 - Number of CPUs 23, 33
 - Number of Nodes 33
 - Number of Sockets 50
 - Number Of Virtual CPUs 46
 - Number of VMs 34
 - Object Name 36
 - Object Status 36
 - Object Type 36
 - Over Under Allocated 28
 - Over Under Allocated CPU 26
 - Over Under Allocated CPU GHz 26
 - Over Under Used 29
 - overview 21
 - Percent CPU Allocated to VMs 27
 - Percent Memory Allocated 29
 - Percent Memory Unallocated 29
 - Percent Memory Unused 30
 - Percent Memory Used 29
 - Percent Used (%) 44
 - Performance Object Status 35
 - Protocol 32
 - Query Name 36
 - Refresh Interval 39
 - Sample Time Delta (sec) 46
 - Sample Timestamp 46
 - Scheduler Parameter Name 41
 - Scheduler Parameter Type 41
 - Scheduler Parameter Value 41
 - Scheduler Parameters 40
 - Sockets Per Node 34
 - Storage Pool Available (GB) 44
 - Storage Pool Capacity (GB) 43
 - Storage Pool Name 42
 - Storage Pool State 43
 - Storage Pool Type 43
 - Storage Pool Used (GB) 43
 - Storage Pool UUID 42
 - Storage Pools 42
 - Threads Per Core 34, 50
 - Timestamp 22, 27, 31, 35, 40, 42, 45
 - Total CPU Time (sec) 47
 - Virtual CPUs Allocated 26
 - Virtual Machine Name 41, 45
 - Virtual Machine State 45
 - Virtual Machine UUID 45
 - Virtual Machines 44
 - Virtualization Type 46
 - VM CPU Fit Estimate 26
 - VM CPU GHz Consumed 24
 - VM CPUs Consumed 23
 - VM Memory Allocated 28

attributes (*continued*)
 VM Memory Allocated Fit Estimate 30
 VM Memory Used 28
 VM Memory Used Fit Estimate 31
Average Collection Duration attribute 39

C

Cache Hit Percent attribute 40
Cache Hits attribute 39
Cache Misses attribute 39
calculate historical data disk space 51
capacity planning for historical data 51
commands
 tacmd addSystem 14
 Take Action 61
components 2
 IBM Tivoli Monitoring 2
configuration 8
 agent 5
 fields 11
 problems and workarounds 77
 remote 14
 values 11
configuring the monitoring agent 5
Cores Per Socket attribute 34, 50
CPU Frequency (GHz) attribute 33
CPU GHz Consumed Per VM attribute 25
CPU GHz Unaccounted For attribute 24
CPU Match attribute 49
CPU Model attribute 32, 49
CPU Percent attribute 47
CPU Percent Consumed by VMs attribute 25
CPU Percent Unaccounted attribute 25
CPU Shares attribute 50
CPU Time Delta (sec) attribute 47
CPUs Consumed attribute 47
CPUs Consumed Per VM attribute 24
CPUs Unaccounted For attribute 24

D

descriptions 54
detailed 71
developerWorks website 98
disk capacity planning for historical data 51
documentation
 See publications

E

enhancements 1
Error Code attribute 37
event
 mapping 91

H

historical data
 calculate disk space 51
 disk capacity planning 51
Host
 situations 55
 workspaces
 descriptions 17

Host CPU attribute group 22
Host CPU Comparison workspace 16
Host Detail workspace 18
Host Memory Allocation Comparison workspace 17
Host Memory attribute group 27
Host Memory Use Comparison workspace 17
Host Name attribute 23, 27, 31, 40, 42, 45
Host Overview workspace 18
Host Relations workspace 18
Host workspace 17
Hosts attribute group 31
Hypervisor URI attribute 32

I

IBM Tivoli Monitoring 2
 overview 1
installation 8
 agent 5
 problems and workarounds 77
 remote 14
installing language packs 5
installing the monitoring agent 5
Integrated Service Management Library documentation 98
interface
 user 3
Intervals Skipped attribute 40
IP Address attribute 32

K

KV1_Host_CPU_Over_Commit_Crit situation 55
KV1_Host_CPU_Over_Commit_Info situation 55
KV1_Host_CPU_Over_Commit_Warn situation 55
KV1_Host_CPU_Pct_High_Crit situation 56
KV1_Host_CPU_Pct_High_Warn situation 56
KV1_Host_Mem_Pct_High_Crit situation 57
KV1_Host_Mem_Pct_High_Warn situation 57
KV1_VM_CPU_Pct_High_Crit situation 58
KV1_VM_CPU_Pct_High_Warn situation 58
KV1_VM_Mem_Pct_High_Crit situation 59
KV1_VM_Mem_Pct_High_Warn situation 58

L

language packs 5
 installing 5
 silent installation 5
Last Collection Duration attribute 39
Last Collection Finished attribute 38
Last Collection Start attribute 38
Linux Kernel-based Virtual Machines
 situations 55
 workspaces
 descriptions 16
Linux Kernel-based Virtual Machines agent
 performance considerations 86
Linux Kernel-based Virtual Machines workspace 16
list of messages 90

M

Memory Allocated (GB) attribute 48
Memory Allocated Per VM attribute 30
Memory Percent attribute 48

- Memory Size (GB) attribute 34
- Memory Size attribute 28
- Memory Used (GB) attribute 48
- Memory Used Per VM attribute 30
- messages
 - contents 89
 - for IBM Tivoli Monitoring for Virtual Environments Agent
 - for Linux Kernel-based Virtual Machines 90
 - format 89
- multi-instance 8
- multiconnection 8

N

- new in this release 1
- Node attribute 22, 27, 31, 35, 40, 42, 44
- Number of Active VMs attribute 25, 30, 35
- Number of Collections attribute 39
- Number of CPU GHz attribute 23, 33
- Number of CPUs attribute 23, 33
- Number of Nodes attribute 33
- Number of Sockets attribute 50
- Number Of Virtual CPUs attribute 46
- Number of VMs attribute 34

O

- Object Name attribute 36
- Object Status attribute 36
- Object Type attribute 36
- operating systems 5
- Over Under Allocated attribute 28
- Over Under Allocated CPU attribute 26
- Over Under Allocated CPU GHz attribute 26
- Over Under Used attribute 29
- overview
 - IBM Tivoli Monitoring 1

P

- Percent CPU Allocated to VMs attribute 27
- Percent Memory Allocated attribute 29
- Percent Memory Unallocated attribute 29
- Percent Memory Unused attribute 30
- Percent Memory Used attribute 29
- Percent Used (%) attribute 44
- performance considerations 86
- Performance Object Status attribute group 35
- policies 63
- prerequisite publications 97
- problems and workarounds 77
 - agent-specific 81
 - agent-specific workspaces 83
 - configuration 77
 - install 77
 - remote deployment 80
 - situations 86
 - Take Action commands 88
 - workspaces 83
- Protocol attribute 32
- publications 97, 98
 - developerWorks website 98
 - IBM Tivoli Monitoring 97
 - Integrated Service Management Library 98
 - prerequisite 97
 - Redbooks 98

- publications (*continued*)
 - related 98
 - Technotes 98
 - wikis 98

Q

- queries, using attributes 21
- Query Name attribute 36

R

- ras1 73
- Redbooks 98
- Refresh Interval attribute 39
- remote
 - installation and configuration 14
- remote deployment
 - problems and workarounds 80
- requirements 5
- response file template 5

S

- Sample Time Delta (sec) attribute 46
- Sample Timestamp attribute 46
- Scheduler Parameter Name attribute 41
- Scheduler Parameter Type attribute 41
- Scheduler Parameter Value attribute 41
- Scheduler Parameters attribute group 40
- silent installation 5
- silent installation of language packs 5
- situations 54
 - additional information
 - predefined, defined 53
 - KV1_Host_CPU_Over_Commit_Crit 55
 - KV1_Host_CPU_Over_Commit_Info 55
 - KV1_Host_CPU_Over_Commit_Warn 55
 - KV1_Host_CPU_Pct_High_Crit 56
 - KV1_Host_CPU_Pct_High_Warn 56
 - KV1_Host_Mem_Pct_High_Crit 57
 - KV1_Host_Mem_Pct_High_Warn 57
 - KV1_VM_CPU_Pct_High_Crit 58
 - KV1_VM_CPU_Pct_High_Warn 58
 - KV1_VM_Mem_Pct_High_Crit 59
 - KV1_VM_Mem_Pct_High_Warn 58
 - overview 53
 - predefined 53
 - problems and workarounds 86
 - Situation Editor 53
- situations, using attributes 21
- Sockets Per Node attribute 34
- SSH protocol 9
- Storage Pool
 - situations 58
 - workspaces
 - descriptions 18
- Storage Pool Available (GB) attribute 44
- Storage Pool Capacity (GB) attribute 43
- Storage Pool Detail workspace 19
- Storage Pool Name attribute 42
- Storage Pool Overview workspace 19
- Storage Pool State attribute 43
- Storage Pool Type attribute 43
- Storage Pool Used (GB) attribute 43
- Storage Pool UUID attribute 42

- Storage Pool workspace 18
- Storage Pools attribute group 42
- support
 - list of messages 90

T

- tacmd addSystem command 14
- Take Action commands
 - additional information 61
 - overview 61
 - predefined 61, 63
 - problems and workarounds 88
- TCP protocol 10
- Technotes 98
- Threads Per Core attribute 34, 50
- Timestamp attribute 22, 27, 31, 35, 40, 42, 45
- Tivoli Enterprise Console
 - event mapping 91
- TLS protocol 10
- Total CPU Time (sec) attribute 47
- trace
 - turn off 76
 - turn on 75
- trace settings 73
- tracing 71
- troubleshooting 65
 - agent-specific 81
 - agent-specific workspaces 83
 - installation 77
 - problems and workarounds 77
 - remote deployment 80
 - situations 86
 - Take Action commands 88
 - turn off trace 76
 - turn on trace 75
 - uninstallation 77
 - workspaces 83

U

- user interface options 3

V

- views
 - Host CPU Comparison workspace 16
 - Host Detail workspace 18
 - Host Memory Allocation Comparison workspace 17
 - Host Memory Use Comparison workspace 17
 - Host Overview workspace 18
 - Host Relations workspace 18
 - Host workspace 17
 - Linux Kernel-based Virtual Machines workspace 16
 - Storage Pool Detail workspace 19
 - Storage Pool Overview workspace 19
 - Storage Pool workspace 18
 - Virtual Machine Detail workspace 19
 - Virtual Machine Overview workspace 20
 - Virtual Machine workspace 19
- Virtual CPUs Allocated attribute 26
- Virtual Machine
 - situations 58
 - workspaces
 - descriptions 19
- Virtual Machine Detail workspace 19

- Virtual Machine Name attribute 41, 45
- Virtual Machine Overview workspace 20
- Virtual Machine State attribute 45
- Virtual Machine UUID attribute 45
- Virtual Machine workspace 19
- Virtual Machines attribute group 44
- virtualization hosts 8
- Virtualization Type attribute 46
- VM CPU Fit Estimate attribute 26
- VM CPU GHz Consumed attribute 24
- VM CPUs Consumed attribute 23
- VM Memory Allocated attribute 28
- VM Memory Allocated Fit Estimate attribute 30
- VM Memory Used attribute 28
- VM Memory Used Fit Estimate attribute 31

W

- wikis 98
- workarounds 77
- Workflow Editor 63
- workspaces
 - descriptions 16
 - Host 17
 - Host CPU Comparison 16
 - Host Detail 18
 - Host Memory Allocation Comparison 17
 - Host Memory Use Comparison 17
 - Host Overview 18
 - Host Relations 18
 - Linux Kernel-based Virtual Machines 16
 - predefined 15
 - problems and workarounds 83
 - Storage Pool 18
 - Storage Pool Detail 19
 - Storage Pool Overview 19
 - Virtual Machine 19
 - Virtual Machine Detail 19
 - Virtual Machine Overview 20
- Workspaces
 - additional information 15
 - overview 15



Printed in USA

SC14-7490-01

