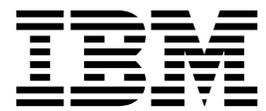


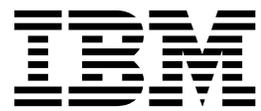
IBM Monitoring Agent for OpenStack
Version 7.5.0.1

User's Guide



IBM Monitoring Agent for OpenStack
Version 7.5.0.1

User's Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 119.

This edition applies to version 7.5.0.1 of IBM Monitoring Agent for OpenStack (product number 5724-L92) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2015.**

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

Contents

Figures	v
--------------------------	----------

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

Chapter 1. Overview of the agent 1

IBM Tivoli Monitoring	1
Functions of the monitoring agent	1
Components of the IBM Tivoli Monitoring environment	1
Agent Management Services	3
User interface options	3

Chapter 2. Agent installation and configuration 5

Prerequisites	5
Language pack installation	6
Installing language packs on Windows systems.	6
Installing language packs on UNIX or Linux systems	6
Silent installation of language packs on Windows, UNIX, or Linux systems	7
Agent-specific installation and configuration.	9
Installing the Monitoring Agent for OpenStack	9
Configuring the agent	16
Configuration values	20
Remote installation and configuration	21

Chapter 3. Workspaces reference 23

Predefined workspaces	23
Workspace descriptions	24
OpenStack Navigator item	24

Chapter 4. Attributes reference 29

Attribute groups for the monitoring agent	29
Attributes in each attribute group	31
API Endpoint Details attribute group.	31
API Endpoints attribute group	33
Environment Summary attribute group	34
Flavors attribute group	37
Hypervisors attribute group	38
Object Containers attribute group	42
Object Storage Account Statistics attribute group	43
Objects attribute group	44
Performance Object Status attribute group	46
Process Details attribute group	49
Processes attribute group	51
Processes by Host attribute group	52
Regions attribute group	53
Take Action Status attribute group.	56
Thread Pool Status attribute group	58

VM Details attribute group	61
Volumes attribute group	64
Disk capacity planning for historical data	66

Chapter 5. Situations reference 69

Predefined situations	69
Situation descriptions	70
OpenStack Navigator item	71

Chapter 6. Take Action commands reference. 73

Predefined Take Action commands	73
---	----

Chapter 7. Policies reference 75

Predefined policies	75
-------------------------------	----

Chapter 8. Troubleshooting 77

Trace logging.	78
Overview of log file management	78
Principal trace log files	79
Examples: Using trace logs	82
RAS trace parameters	83
Dynamic modification of trace settings	85
Setting trace parameters for the Tivoli Enterprise Console server	88
Problems and workarounds	88
Installation and configuration troubleshooting.	89
Remote deployment troubleshooting	92
Agent troubleshooting.	93
Workspace troubleshooting	97
Situation troubleshooting	100
Take Action commands troubleshooting	102
Support information	102
Informational, warning, and error messages overview	103
Message format	103
Agent messages	104

Appendix A. Event mapping 105

Appendix B. Documentation library 115

Prerequisite publications.	115
Related publications	116
Other sources of documentation	116

Notices 119

Trademarks	121
----------------------	-----

Index 123

Figures

Tables

1. Silent installation parameters	11	6. General problems and solutions for uninstallation	91
2. Capacity planning for historical data logged by the Monitoring Agent for OpenStack	67	7. Remote deployment problems and solutions	92
3. Information to gather before contacting IBM Software Support	77	8. Agent problems and solutions	93
4. Trace log files for troubleshooting agents	79	9. Workspace problems and solutions.	97
5. Problems and solutions for installation and configuration	89	10. Situation problems and solutions	100
		11. Take Action commands problems and solutions	102

Chapter 1. Overview of the agent

The IBM® Monitoring Agent for OpenStack (product code SG) provides you with a central point of management for your OpenStack environment or application.

By using the Monitoring Agent for OpenStack you can collect and analyze complex OpenStack environment in an easy way. The agent alerts you when an event is triggered and help you to determine which part of the system has an error. The Monitoring Agent for OpenStack monitors related resource in the following layers:

- OpenStack infrastructure and services run time, for example, processes status and API health status.
- Cloud resources that are controlled by OpenStack, including running cloud instance amount, total instance amount, total CPU or core number, total memory size (GB), total disk (number), total images, and resources used percentage.
- Lifecycle of workload (Virtual Server) instances that are administered by OpenStack, for example, workload instance status and workload instance metrics available (including CPU, Memory, and Network usage).

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 115 for complete information about IBM Tivoli Monitoring and the Tivoli Enterprise Portal.

Functions of the monitoring agent

Components of the IBM Tivoli Monitoring environment

After you install and set up the Monitoring Agent for OpenStack, 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, Monitoring Agent for OpenStack

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 Tivoli Integrated Portal.

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.

Tivoli Integrated Portal

Tivoli Integrated Portal helps the interaction and secure passing of data between Tivoli products through a common portal. You can launch from one application to another and within the same

dashboard view research different aspects of your managed enterprise. This component is installed automatically with the first Tivoli product that uses the Tivoli Integrated Portal framework. Subsequent products can install updated versions of Tivoli Integrated Portal. After version 2.2, this component is replaced by the Dashboard Application Services Hub.

Agent Management Services

You can use IBM Tivoli Monitoring Agent Management Services to manage the Monitoring Agent for OpenStack.

Agent Management Services is available for the following IBM Tivoli Monitoring OS agents: Windows, Linux, and UNIX. The services are designed to keep the Monitoring Agent for OpenStack 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 for Tivoli Common Reporting 3.1 and on Tivoli Integrated Portal for earlier versions.

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.

Tivoli Integrated Portal

Web-based products that are built on the Tivoli Integrated Portal framework share a common user interface where you can launch applications and share information. After version 2.2, this interface is replaced by the Dashboard Application Services Hub.

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 Monitoring Agent for OpenStack, 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*.

With the self-describing agent capability, new or updated IBM Tivoli Monitoring agents using IBM Tivoli Monitoring V6.2.3 or later can become operational after installation without having to perform additional product support installation steps. To take advantage of this capability, see “Enabling self-describing agent capability at the hub monitoring server” in the *IBM Tivoli Monitoring Installation and Setup Guide*. Also, see “Self-describing monitoring agents” in the *IBM Tivoli Monitoring Administrator’s Guide*.

Prerequisites

The IBM Monitoring Agent for OpenStack uses the OpenStack Python API (component client) to collect monitored resource metrics remotely. Corresponding Python run-time libraries and OpenStack clients are required. The agent also requires the administrator authorization to access the OpenStack components through keystone service. If you want to collect OpenStack process-related information, you need the credentials to access the remote servers through ssh.

Software prerequisites

The operating systems that are supported by IBM Monitoring Agent for OpenStack are 64-bit Redhat 6.5, 64-bit Redhat 7, and SuSE 11.

The Monitoring Agent for OpenStack can work with IBM Tivoli Monitoring 6.2.3 FP5, and IBM Tivoli Monitoring 6.3.0 or later.

Before you can use the Monitoring Agent for OpenStack, you must have the following software on the server where you install the agent:

- Python 2.6.0 or later, or Python 2.7.0 or later.
- OpenStack clients:
 - Keystone
 - Nova
 - Neutron
 - Cinder
 - Swift
- Supported OpenStack client versions:
 - Kilo
 - Juno
 - IceHouse

To install the OpenStack command-line clients, refer to the topic *Install the OpenStack command-line clients* (http://docs.openstack.org/cli-reference/content/install_clients.html).

- Paramiko library for remote access in Python.
- KornShell.

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”
- “Silent installation of language packs on Windows, UNIX, or Linux systems” on page 7

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 lpinstaller.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 `n1spackage` 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 8.
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 `n1spackage` folder, for example:
`NLS_PACKAGE_FOLDER = //tmp//LP//n1spackage.`

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,
#//installtiviali//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//nlspackge//kex_nls.nlspkg;//  
tmp//LP//nlspackge//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 Monitoring Agent for OpenStack.

Installing the Monitoring Agent for OpenStack

Installation involves installing an agent and installing agent application support on IBM Tivoli Monitoring components.

Procedure

Complete the following steps to install the Monitoring Agent for OpenStack:

1. Decide whether you want to use the self-describing capability. For detailed information about the self-describing capability, see “Enabling the self-describing capability to automatically install application support” on page 13.
2. Install the Monitoring Agent for OpenStack. For instructions, see “Installing an agent.”
3. (Self-describing capability is disabled) Install the application support on Tivoli Enterprise Monitoring Server, Tivoli Enterprise Portal Server, and Tivoli Enterprise Portal desktop client. For instructions, see “Manually installing application support” on page 14.
4. (Self-describing capability is enabled) If you are using the Tivoli Enterprise Portal desktop client, install the application support on the portal desktop client. For instructions, see “Installing Tivoli Enterprise Portal desktop client support” on page 15.

Installing an agent

You must install the Monitoring Agent for OpenStack before you can monitor your OpenStack environment. The agent can be installed through a refresh installation only.

Before you begin

Make sure that the software listed in “Prerequisites” on page 5 is installed.

If you want to enable self-describing agent capability, see “Enabling the self-describing capability to automatically install application support” on page 13

About this task

To avoid high CPU and memory load on the server, you can install the Monitoring Agent for OpenStack on a stand-alone server and remotely monitor the OpenStack environment.

Procedure

You have two options to install the Monitoring Agent for OpenStack on a Linux system:

- To install the agent by running the script and responding to prompts, see “Installing the agent by responding to prompts” on page 10.
- To install the agent by editing the silent response file and running the script with no interaction, see “Performing a silent installation” on page 10.

Installing the agent by responding to prompts:

Procedure

1. Download and extract the agent installation image to the system where you want to install the Monitoring Agent for OpenStack.
2. In the directory where you extracted the agent installation image, run the following command:
`./install.sh`
3. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory `/opt/IBM/ITM` or type a different directory where you want to install the agent.
4. If the installation directory does not exist, you are asked if you want to create it. Type `y` to create this directory and press Enter.
5. If any existing IBM Tivoli Monitoring components are currently running on the computer, the installer stops the components during the installation process, then restarts the components when the installation is complete. To confirm this action, type `1` when prompted and press Enter. If you choose not to stop the components, the installation process ends.
6. When prompted to install products to the local host, type `1` and press Enter.
7. When the software license agreement is displayed, type `1` to accept the agreement and press Enter.
8. If IBM GSKit is not installed on the computer, you are prompted to provide an encryption key. Use the same key across the enterprise. Either type the key or accept the default and press Enter.
9. When a numbered list of available operating systems is displayed, type `1` to install the IBM Tivoli Monitoring components for this operating system and press Enter.
10. When a numbered list of available components is displayed, type the number corresponding to IBM Monitoring Agent for OpenStack and press Enter.
11. Type `1` to confirm your selection.
12. At the prompt Do you want to install additional products or product support packages, type `2` and press Enter.
13. If your IBM Tivoli Monitoring environment is not secured, you are asked if you want to secure it. If your IBM Tivoli Monitoring environment is already secured, this question is skipped. The product installation process creates most of directories and files with write permissions. You can secure your installation or choose the default value.

Results

The Monitoring Agent for OpenStack is installed.

What to do next

If you enabled self-describing feature for the Monitoring Agent for OpenStack, the agent will automatically push the application support files to Tivoli Enterprise Monitoring Server and Tivoli Enterprise Portal Server. You only need to manually install the application support files on Tivoli Enterprise Portal desktop client if you are using the desktop client. For instructions on how to install application support on the portal desktop client, see “Installing Tivoli Enterprise Portal desktop client support” on page 15.

If you do not use the self-describing feature, you must manually the application support on IBM Tivoli Monitoring components. For detailed information, see “Manually installing application support” on page 14.

Performing a silent installation:

About this task

Silent installation uses a response file to define the installation parameters for the Monitoring Agent for OpenStack. You can edit the parameters in this response file to control the installation options. A sample response file is shipped with the agent installation image. The file name is `silent_install.txt`.

Procedure

1. Download and extract the installation image to the system where you want to install the agent.
2. In the root directory of the extracted installation image, locate the `silent_install.txt` file.
3. Edit the `silent_install.txt` file to set the parameters described in Table 1 to appropriate values for your environment and save the changes.

Table 1. Silent installation parameters

Parameter	Definition
INSTALL_ENCRYPTION_KEY	<p>REQUIRED. The data encryption key that is used to encrypt data that is sent between systems. This key must be the same for all components in your IBM Tivoli Monitoring environment.</p> <p>Do not use the following characters in the key:</p> <ul style="list-style-type: none"> & ampersand pipe ' single quote = equal sign \$ dollar sign <p>In addition, do not specify double-byte (DBCS) characters.</p>
INSTALL_FOR_PLATFORM	<p>The operating system for which to install the components. You can specify an architecture code. If you do not specify an architecture code, the operating system for the current computer is used. You can find a list of the architecture codes for the supported architectures in <code>archdsc.tbl</code> in the registry directory.</p>
INSTALL_PRODUCT	<p>The product code for the components (or "products") that you want to install.</p> <p>Note: The product code for the Monitoring Agent for OpenStack is sg.</p> <p>You can use the <code>./cinfo</code> command to view the product codes for the applications that are installed on this computer. You can also find a list of the product codes in the registry directory in <code>proddsc.tbl</code>.</p> <p>You can specify "all" to install all available components.</p> <p>To install multiple components (but not all), repeat this parameter for each component that you want to install. For example:</p> <pre>INSTALL_PRODUCT=sg INSTALL_PRODUCT=cj INSTALL_PRODUCT=cq</pre> <p>This example installs the monitoring server, portal server, and portal desktop client on a Linux computer.</p> <p>In the same pass, you can install support with the components. After you specify the components that you want to install, specify the support that you want to install by using the following parameters:</p> <pre>INSTALL_PRODUCT_TMS (for TEMS support) INSTALL_PRODUCT_TPS (for TEPS support) INSTALL_PRODUCT_TPW (for TEP Browser Client support) INSTALL_PRODUCT_TPD (for TEP Desktop Client support) INSTALL_PRODUCT_TPA (for ITPA Domain support)</pre>

Table 1. Silent installation parameters (continued)

Parameter	Definition
MS_CMS_NAME	If you are installing a monitoring server, use this parameter to specify the name for the monitoring server, such as HUB <i>hostname</i> . Do not specify an IP address or fully qualified host name.
SKIP_SDA_CHECK	You can use this parameter to overwrite product support that was seeded in self-describing mode. By default, this parameter is set to NO and a self-describing mode seeding status check is performed. If this parameter is set to YES , the seeding process at the end of the installation does not check the self-describing mode seeding status.
PREREQ_CHECK	Set this parameter to PREREQ_CHECK=YES to identify any missing system prerequisites before you begin an installation or upgrade to a new version of IBM Tivoli Monitoring.

4. Run the following command to install Monitoring Agent for OpenStack:

```
./install.sh -q -h install_dir -p response_file
```

Where :

install_dir

Specifies the installation location for the Monitoring Agent for OpenStack. The default installation location is /opt/IBM/ITM.

Do not use the number sign (#) or exclamation mark (!) characters anywhere in the specified installation path.

Important: You must not specify the path of the directory that contains the `./install.sh` script as your installation directory. On certain operating systems, this can cause the plug-in JAR files to overwrite themselves and become zero length files. The installation will fail as a result.

response_file

Specifies the full path of the `silent_install.txt` response file that you use to specify the installation parameters.

The following parameters are available for the `./install.sh` script:

[-k] Optional.

Secures your IBM Tivoli Monitoring installation. If you do not secure the installation at this point you will be asked at the end of this procedure if you want to do so.

[-c] Optional.

Prints diagnostic messages to the console

[-d] Optional.

Indicates the location of the product installation image.

[-h] Optional.

Indicates the installation directory of IBM Tivoli Monitoring components.

[-j] Optional.

Indicate the JRE installed location. This parameter is available on Tandem system only.

[-p] Optional.

Specifies the silent installation response file with installation parameters specified.

Results

The Monitoring Agent for OpenStack is installed.

What to do next

If you enabled self-describing feature for the Monitoring Agent for OpenStack, the agent will automatically push the application support files to Tivoli Enterprise Monitoring Server and Tivoli Enterprise Portal Server. You only need to manually install the application support files on Tivoli Enterprise Portal desktop client if you are using the desktop client. For instructions on how to install application support on the portal desktop client, see “Installing Tivoli Enterprise Portal desktop client support” on page 15.

If you do not use the self-describing feature, you must manually the application support on IBM Tivoli Monitoring components. For detailed information, see “Manually installing application support” on page 14.

Enabling the self-describing capability to automatically install application support

The Monitoring Agent for OpenStack supports the self-describing capability, which can automatically install agent application support on IBM Tivoli Monitoring servers. Before you install a self-describing agent, decide whether you want to enable the self-describing capability.

About this task

By default, the self-describing environment variable is set to Y to enable the self-describing capability on the following IBM Tivoli Monitoring components:

- Any remote Tivoli Enterprise Monitoring Server the agent is connected to (**KMS_SDA**)
- Tivoli Enterprise Portal Server (**TEPS_SDA**)
- Monitoring agent that supports the self-describing capability (**TEMA_SDA**)

You can disable the self-describing capability by changing the self-describing environment to N for each component.

A best practice is to accept the default settings and use the self-describing environment variable only at the hub monitoring server to control the self-describing agent enablement.

Procedure

To enable and use the self-describing capability to automatically install agent application support, complete the following steps:

1. Configure the hub monitoring server to set the **KMS_SDA** variable to Y.
2. (IBM Tivoli Monitoring V6.3 or later) Run the following command to permit self-describing agent installations without any blocking:

```
tacmd editsdainstalloptions -t DEFAULT -i ON
```

Tip: The command permits the application support installation for all self-describing agents. To permit self-describing agent installation only for the Monitoring Agent for OpenStack, use the following command:

```
tacmd addsdainstalloptions -t sg -v 01000000
```

For more information about the self-describing capability, see *Dynamically controlling the hub monitoring server self-describing agent capability* section in the *IBM Tivoli Monitoring: Installation and Setup Guide*.

Manually installing application support

You must install application support before you can view data collected by the Monitoring Agent for OpenStack.

Procedure

Follow the steps in the following topics to install application support manually:

- Installing Tivoli Enterprise Monitoring Server support.
- Installing Tivoli Enterprise Portal Server support.
- Installing Tivoli Enterprise Portal Desktop Client support.

Installing Tivoli Enterprise Monitoring Server support:

Install the application support for the Monitoring Agent for OpenStack to work properly with Tivoli Enterprise Monitoring Server.

Procedure

Complete the following steps to install Tivoli Enterprise Monitoring Server support:

1. In the directory where you extracted the installation files, run the following command:

```
./install.sh
```
2. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory `/opt/IBM/ITM` or type the full path of the monitoring server installation directory.
3. If a message is displayed stating that the directory already exists and asking whether to use it, type 1 and press Enter.
4. If any existing IBM Tivoli Monitoring components are currently running on the computer, the installation program stops the components during the installation process, then restarts the components when the installation is complete. To confirm this action, type 1 when prompted and press Enter. If you choose not to stop the components, the installation process ends.
5. When prompted to install products to the local host, type 1 and press Enter.
6. When the software license agreement is displayed, type 1 to accept the agreement and press Enter.
7. If IBM GSKit is not installed on the computer, you are prompted to provide an encryption key. Use the same key across the enterprise. Either type the key or accept the default and press Enter.
8. When a numbered list of available components is displayed, type the number for Tivoli Enterprise Monitoring Server support and press Enter.
9. Type 1 to confirm your selection and press Enter.
10. When a list of products for which the support files are to be added is displayed, type the number corresponding to IBM Monitoring Agent for OpenStack and press Enter.
11. Type 1 to confirm your selection and press Enter.
12. At the prompt `Do you want to install additional products or product support packages`, type 2 and press Enter.
13. At the prompt `Do you want to seed product support on the Tivoli Enterprise Monitoring Server`, type 1 and press Enter.
14. When prompted to select support packages to be seeded, type 1 and press Enter.
15. If your IBM Tivoli Monitoring environment is not secured, you are asked if you want to secure it. If your IBM Tivoli Monitoring environment is already secured, this question is skipped. The product installation process creates most of directories and files with write permissions. You can secure your installation or choose the default value.

Results

Tivoli Enterprise Monitoring Server support is installed.

Installing Tivoli Enterprise Portal Server support:

Install the application support for the Monitoring Agent for OpenStack to work properly with Tivoli Enterprise Portal Server.

Procedure

Complete the following steps to install Tivoli Enterprise Portal Server support:

1. In the directory where you extracted the installation files, run the following command:
`./install.sh`
2. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory `/opt/IBM/ITM` or type the full path of the portal server installation directory.
3. If a message is displayed stating that the directory already exists and asking whether to use it, type 1 and press Enter.
4. If any existing IBM Tivoli Monitoring components are currently running on the computer, the installer stops the components during the installation process, then restarts the components when the installation is complete. To confirm this action, type 1 when prompted and press Enter. If you choose not to stop the components, the installation process ends.
5. When prompted to install products to the local host, type 1 and press Enter.
6. When the software license agreement is displayed, type 1 to accept the agreement and press Enter.
7. If IBM GSKit is not installed on the computer, you are prompted to provide an encryption key. Use the same key across the enterprise. Either type the key or accept the default and press Enter.
8. When a numbered list of available components is displayed, type the number for Tivoli Enterprise Portal Server support and press Enter.
9. Type 1 to confirm your selection and press Enter.
10. When a list of products for which the support files are to be added is displayed, type the number corresponding to IBM Monitoring Agent for OpenStack and press Enter.
11. Type 1 to confirm your selection and press Enter.
12. If the system is running Tivoli Enterprise Portal Server and you want to install Tivoli Enterprise Portal Browser Client support, at the prompt Do you want to install additional products or product support packages, type 1 and press Enter.
13. At the list of component support categories, type the number for Tivoli Enterprise Portal Browser Client support and press Enter.
14. Repeat from step 7 through step 11, selecting Tivoli Enterprise Portal Browser Client Support.
15. At the prompt Do you want to install additional products or product support packages, type 2 and press Enter.

Results

Tivoli Enterprise Portal Server support is installed.

Installing Tivoli Enterprise Portal desktop client support:

Install the application support for the Monitoring Agent for OpenStack to work properly with Tivoli Enterprise Portal desktop client.

Procedure

Complete the following steps to install Tivoli Enterprise Portal desktop client support:

1. In the directory where you extracted the installation files, run the following command:
`./install.sh`
2. When prompted for the IBM Tivoli Monitoring home directory, press Enter to accept the default directory `/opt/IBM/ITM` or type the full path of the portal desktop client installation directory.
3. If a message is displayed stating that the directory already exists and asking whether to use it, type 1 and press Enter.
4. If any existing IBM Tivoli Monitoring components are currently running on the computer, the installer stops the components during the installation process, then restarts the components when the installation is complete. To confirm this action, type 1 when prompted and press Enter. If you choose not to stop the components, the installation process ends.
5. When prompted to install products to the local host, type 1 and press Enter.
6. When the software license agreement is displayed, type 1 to accept the agreement and press Enter.
7. If IBM GSKit is not installed on the computer, you are prompted to provide an encryption key. Use the same key across the enterprise. Either type the key or accept the default and press Enter.
8. When a numbered list of available components is displayed, type the number for Tivoli Enterprise Portal desktop client support and press Enter.
9. Type 1 to confirm your selection and press Enter.
10. When a list of products for which the support files are to be added is displayed, type the number corresponding to IBM Monitoring Agent for OpenStack and press Enter.
11. Type 1 to confirm your selection and press Enter.
12. At the prompt Do you want to install additional products or product support packages, type 2 and press Enter.

Results

Tivoli Enterprise Portal desktop client support is installed.

Configuring the agent

You must configure the Monitoring Agent for OpenStack before the agent can automatically monitor the OpenStack environment.

Procedure

1. You can either configure the agent by responding to prompts or by using a response file. For instructions, see “Configuring the agent.”
2. If you want to collect process-related information, configure the data collector for the Monitoring Agent for OpenStack. For instructions, see Configuring agent data collector and SSH connections.

Configuring the agent

For a typical environment, if you want the Monitoring Agent for OpenStack to automatically monitor the OpenStack environment, you must configure the agent first.

Procedure

You have two options to configure the Monitoring Agent for OpenStack on a Linux system:

- To configure the agent by running the script and responding to prompts, see “Configuring the agent by responding the prompts” on page 17.
- To configure the agent by editing the silent response file and running the script with no interaction, see “Configuring the agent with a response file” on page 18.

Configuring the agent by responding the prompts:

Procedure

1. To configure the agent, run the following command:

```
install_dir/bin/itmcmd agent config -A sg
```

Where *install_dir* is the IBM Tivoli Monitoring installation directory. The default IBM Tivoli Monitoring installation directory is `/opt/IBM/ITM`.

2. When prompted to Enter instance name, input an instance name and enter 1 to continue.

Important: The Monitoring Agent for OpenStack is a multiple instance agent and requires an instance name for each agent instance. The instance name that you specify is included in the managed system name `instance_name:host_name:sg`. The length of the instance name you specify is limited to 28 characters minus the length of your host name. For example, if you specify `OS1` as your instance name, your managed system name is `OS1:hostname:SG`.

3. When prompted to Edit Monitoring Agent for OpenStack, enter 1 to continue.
4. When prompted to Edit OpenStack environment authentication information, enter 1 to continue and provide the following information:

```
OpenStack authentication url (default is: http://localhost:5000/v2.0):
OpenStack username (default is: admin):
Enter OpenStack password (default is: ):
Re-type: OpenStack password (default is: ):
OpenStack tenant name (default is: admin):
```
5. When prompted to Edit Python, enter 1 to continue.
6. When prompted for Python Executable Location, specify the Python executable location, for example, `/usr/bin/python`. You can find the fully qualified path by running the following command in your environment:

```
which python
```
7. When prompted to Edit Socket settings, enter 1 to continue.
8. When prompted for Port Number, accept the default value or specify a port number. This port is used for monitoring internal communication between the Monitoring Agent for OpenStack data collector and the Monitoring Agent for OpenStack, both of which are installed on a local server only. The agent listens on this port for data from the data collector. The default value of 0 indicates that an ephemeral port is used when the agent starts. On a server with strict security rules on ports, you can configure one specific port for the agent to use. This port is for internal use by the agent and is not related to OpenStack environment.
9. Press Enter when you are asked if the agent connects to a Tivoli Enterprise Monitoring Server.
10. Type the protocol that you want to use to communicate with the Tivoli Enterprise Monitoring Server.
11. If you want to set up a backup protocol, enter that protocol and press Enter.
12. When prompted for TEMS Host Name for IPv4, type the name of the host system on which the Tivoli Enterprise Monitoring Server that you want the agent to connect to is installed and press Enter.
13. Depending on the type of protocol you specified, provide the port number and press Enter.
14. When you are asked to Configure connection for a secondary TEMS, provide the secondary TEMS information if you have a secondary TEMS.
15. After you finish configuring the agent, you can start the agent instance by running the command:

```
install_dir/bin/itmcmd agent -o instance_name start sg
```

where *instance_name* is the name of the agent instance you configure.

Results

The agent is configured successfully.

What to do next

If you want to collect process-related information, configure the data collector for the Monitoring Agent for OpenStack by completing the steps in *Configuring agent data collector and SSH connections*.

If you want to change the trace level of the agent for troubleshooting purpose, edit the value of the variable **KBB_RAS1** in the *install_dir/config/sg.ini* file according to the instruction in the file.

For more information about the agent environment variables, see the *Environment variables* topics in *Appendix E. Agent configuration and environment variables* of the *IBM Tivoli Monitoring: Installation and Setup Guide*.

Configuring the agent with a response file:

Procedure

1. Go to the *install_dir/sample* directory and edit the *sg_silent_config.txt* configuration file for the Monitoring Agent for OpenStack.
2. Specify values for the parameters that are identified in the file. The response file contains comments that define the available parameters and the values to specify.
3. Save the file and exit.
4. Run the following command to configure the agent:

```
./itmcmd config -A -p response_file sg
```

where *response_file* is the name of the configuration response file. Specify an absolute path to this file.

Results

The agent is configured.

What to do next

If you want to collect process-related information, configure the data collector for the Monitoring Agent for OpenStack by completing the steps in *Configuring agent data collector and SSH connections*.

If you want to change the trace level of the agent for troubleshooting purpose, edit the value of the variable **KBB_RAS1** in the *install_dir/config/sg.ini* file according to the instruction in the file.

For more information about the agent environment variables, see the *Environment variables* topics in *Appendix E. Agent configuration and environment variables* of the *IBM Tivoli Monitoring: Installation and Setup Guide*.

Configuring agent data collector and SSH connections

If you want to collect process-related information, configure the agent data collector for the Monitoring Agent for OpenStack and set up SSH connections with the target OpenStack component server.

About this task

You must set up SSH connection to collect process information before you start the Monitoring Agent for OpenStack. To do this, you can use the product provided **ksg_setup_key.sh** or **ksg_ssh_setup.py** assistance tool as indicated in the following procedure.

If you are familiar with setting up SSH connections, you can also use the **ssh-keygen** and **ssh-copy-id** Linux commands to do it.

Procedure

1. Go to the `install_dir/config` directory, where `install_dir` is the IBM Tivoli Monitoring installation directory.
2. Edit the `ksg_dc_instance_name.cfg` file, where `instance_name` is the name you specified for this agent instance. The file is created after the agent instance starts. If the file does not exist, copy `install_dir/$ITM_BINARCH/sg/bin/ksg_dc.cfg` to the directory `install_dir/config` and change the file name to `ksg_dc_instance_name.cfg`. For example, if the instance name is `OS1`, change the name to `ksg_dc_OS1.cfg`.
3. In the `ksg_dc_instance_name.cfg` file, set the value of the parameter **collect_process_information** to `YES`.
4. In the section **OS_process_collection**, specify value for the parameter **ssh_user_host** with the users and host names or IP addresses of the OpenStack component servers according to the format of the following example:

```
ssh_user_host=root@9.112.250.248,user1@hostname
```
5. Save the settings.
6. For the settings to take effect, restart the agent instance by running the following commands:

```
install_dir/bin/itmcmd agent -o instance_name stop sg
install_dir/bin/itmcmd agent -o instance_name start sg
```
7. Set up the SSH connections with the target component server by using one of the following ways:
 - Set up the connections one by one by using the `ksg_setup_key.sh` script . To do that, go to `install_dir/$ITM_BINARCH/sg/bin` directory and run the `ksg_setup_key.sh` script with host name or IP and user to build the SSH connections with component servers that are specified in Step 4. If you follow the example that is given in Step 4, you must run the script twice to set up the connection one by one:

```
./ksg_setup_key.sh 9.112.250.248 root
./ksg_setup_key.sh hostname user1
```

Note: You must provide the passwords when you run the scripts for the first time. You don't need to provide the passwords again.

- Set up the connections one by one or in a batch job by using the `ksg_ssh_setup.py` tool that is provided by Monitoring Agent for OpenStack in `install_dir/$ITM_BINARCH/sg/bin`. You must install Python `pepexpect` library before you can use this tool.
 - To set up SSH connections one by one, run the command:

```
python ksg_ssh_setup.py -single
```

This command helps you to set up the SSH connection to the remote target server. You must provide the following information.

```
Enter the remote target machine host name or IP address: (Type 'END' to end input.)
Enter the account to access the remote machine(e.g. root):
Enter the above user's password:
```

- To set up SSH connections in a batch job, run the command:

```
python ksg_ssh_setup.py -ssh SSH_file
```

where `SSH_file` is the file that contains the target server, user, and password information. You must create the file according to the `ksg_dc_ssh_list.txt` file in the same directory as the Python tool, and specify the host and user information in the file according to the format of the examples:

```
hostname root passw0rd
9.112.250.248 user1 passw0rd
```

Note: No need to set up the connections again after you restart the agent or change the agent configuration. You must set up the connections again only when the user name or password for the target server change.

Results

The data collector is configured and the SSH connections are set up properly.

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 that you type are not displayed to help maintain the security of these values.

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

OpenStack environment authentication information (OPENSTACK_CONNECTION)

The OpenStack environment authentication information

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

This group defines information that applies to the entire agent.

OpenStack authentication url (KSG_OPENSTACK_AUTH_URL)

The auth_url of OpenStack environment

The type is string.

This value is required.

Default value: http://localhost:5000/v2.0

Note: Monitoring Agent for OpenStack supports only OpenStack Identity API v2 through HTTP.

OpenStack password (KSG_OPENSTACK_PASSWORD)

The administrator user's password

The type is password.

This value is required.

Default value: None

OpenStack tenant name (KSG_OPENSTACK_TENANT_NAME)

The OpenStack tenant name, also known as project name

The type is string.

This value is required.

Default value: admin

OpenStack username (KSG_OPENSTACK_USERNAME)

The administrator user to login the OpenStack environment

The type is string.

This value is required.

Default value: admin

Python (KSG_PYTHON)

Python executable location

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

This group defines information that applies to the entire agent.

Python Executable Location (KSG_PYTHON_LOCATION)

The python executable which will be used to run OpenStack agent data collector. You can find the fully qualified path by running the following command in your terminal: "which python".

The type is string.

This value is required.

Default value: None

Socket (KSG_SOCKET)

Socket Data Source

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

This group defines information that applies to the entire agent.

Port Number (CP_PORT)

The port that the agent will use to listen on for data from socket clients. A value of 0 indicates an ephemeral port will be used. Please note that this does NOT correspond to any ports used by your application(s). This port is for internal use by the agent.

The type is numeric.

This value is optional.

Default value: 0

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 20.

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 SG -n Primary:sample.node.name:NT
-p KSG_PYTHON.KSG_PYTHON_LOCATION=value
  KSG_SOCKET.CP_PORT=value
  OPENSTACK_CONNECTION.KSG_OPENSTACK_AUTH_URL=value
  OPENSTACK_CONNECTION.KSG_OPENSTACK_PASSWORD=value
  OPENSTACK_CONNECTION.KSG_OPENSTACK_TENANT_NAME=value
  OPENSTACK_CONNECTION.KSG_OPENSTACK_USERNAME=value
INSTANCE="inst1"
```

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 29.

Predefined workspaces

The Monitoring Agent for OpenStack provides predefined workspaces, which are organized by Navigator item.

- OpenStack Navigator item
 - OpenStack workspace
 - API Endpoints workspace
 - API Endpoints Details workspace
 - Environment Summary workspace
 - Flavors workspace
 - Hypervisors workspace
 - Object Containers workspace

- Object Storage Account Statistics workspace
- Objects workspace
- Process Details workspace
- Processes workspace
- Regions workspace
- Virtual Machines workspace
- Volumes 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.

OpenStack Navigator item

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

OpenStack workspace

This workspace contains the following view:

API Endpoints workspace

The API Endpoints workspace displays the API endpoints status summarized by OpenStack service type in monitored OpenStack environment. You can link to API Endpoints Details workspace from API Endpoints navigation node to check the detailed information of all the API Endpoints. To link to that workspace, right-click the API Endpoints navigation node and select API Endpoints Details link.

This workspace contains the following views:

API Endpoints

The API Endpoints view shows a table view displaying the number of available or unavailable API endpoints of each OpenStack component. You can link to API Endpoints Details workspace from API Endpoints table view to check the API Endpoints Details of a specific OpenStack component. To link to that workspace, right-click the row that represents the OpenStack component you want to check then select the link to API Endpoint Details in the context menu.

API Endpoint Statistics by Component

The API Endpoint Statistics by Component view displays the number of available or unavailable API endpoints of each OpenStack component. You can link to API Endpoints Details workspace from API Endpoint Statistics by Component view to check the API Endpoints Details of a specific OpenStack component. To link to that workspace, right-click the bar that represents the OpenStack component you want to check and click Link To, then click API Endpoints Details.

Total API Endpoints

The Total API Endpoints view displays the number of total available API endpoints and the number of total unavailable API endpoints.

API Endpoints Details workspace

The API Endpoints Details workspace is a link to workspace that displays the detailed information about each API endpoint, such as its region name, service type, public URL, status, detecting information and http status code.

This workspace contains the following view:

API Endpoints Details

The API Endpoints Details view shows a table view displaying the detailed information about API endpoints.

Environment Summary workspace

The Environment Summary workspace displays the overall status of the monitored OpenStack environment. You can open this workspace by clicking the OpenStack Agent navigation node. You can link to Object Storage Account Statistics workspace from OpenStack Agent navigation node to check the object storage account information for current tenant. To link to that workspace, right-click the OpenStack Agent navigation node and then select the link to Object Storage Account Statistics in the context menu.

This workspace contains the following views:

Environment Summary

The Environment Summary view shows a table view displaying the overall status of the monitored OpenStack environment, such as the number of processes, API endpoints, virtual machine instances in the environment, and the number of regions with different resource usages. You can link to Object Storage Account Statistics workspace from Environment Summary table view to check the object storage account information for current tenant. To link to that workspace, right-click the row that represents OpenStack environment summary and then select the link to Object Storage Account Statistics in the context menu.

API Endpoint Statistics by Component

The API Endpoint Statistics by Component view displays the number of available or unavailable API endpoints in monitored OpenStack environment grouped by OpenStack component.

Process Count by Component

The Process Count by Component view displays the number of running processes in monitored OpenStack environment grouped by OpenStack component.

Flavors workspace

The Flavors workspace is a link to workspace that displays the information about each flavor, such as its ID, name, configurations and reference count.

This workspace contains the following view:

Flavors

The Flavors view shows a table view displaying the detailed information about flavors.

Hypervisors workspace

The Hypervisors workspace displays the hypervisors information and the resource usages of the hypervisors in all the monitored OpenStack regions. You can link to Virtual Machine workspace from Hypervisors navigation node to check the virtual machines on all the hypervisors. To link to that workspace, right-click the Hypervisors navigation node and select Virtual Machines link.

This workspace contains the following views:

Hypervisor Summary

The Hypervisor Summary view shows a table view displaying the hypervisors information and the resource usages of all the hypervisors, such as each hypervisor's VCPU, Memory, File System Storage usage percentages and the number of running virtual machine instances on each hypervisor. You can link to Virtual Machines workspace from Hypervisor Summary table view to check the virtual machines created by a specific hypervisor. To link to that workspace, right-click the row that represents the hypervisor you want to check, then select link to Virtual Machines in the context menu. .

Top 5 Hypervisors by Count of Running VMs

The Top 5 Hypervisors by Count of Running VMs view displays the top 5 hypervisors that have the most running virtual machines. You can link to Virtual Machines workspace from Top 5 Hypervisors by Count of Running VMs view to check the virtual machines created by a specific hypervisor. To link to that workspace, right-click the bar that represents the hypervisor you want to check and click Link To, then click Virtual Machines.

Top 5 Hypervisors by VCPU Usage Percentage

The Top 5 Hypervisors by VCPU Usage Percentage view displays the top 5 hypervisors that have the highest VCPU usage percentage. You can link to Virtual Machines

workspace from Top 5 Hypervisors by VCPU Usage Percentage view to check the virtual machines created by a specific hypervisor. To link to that workspace, right-click the bar that represents the hypervisor you want to check and click Link To, then click Virtual Machines.

Top 5 Hypervisors by Memory Usage Percentage

The Top 5 Hypervisors by Memory Usage Percentage view displays the top 5 hypervisors that have the highest memory usage percentage. You can link to Virtual Machines workspace from Top 5 Hypervisors by Memory Usage Percentage view to check the virtual machines created by a specific hypervisor. To link to that workspace, right-click the bar that represents the hypervisor you want to check and click Link To, then click Virtual Machines.

Top 5 Hypervisors by Storage Usage Percentage

The Top 5 Hypervisors by Storage Usage Percentage view displays the top 5 hypervisors that have the highest file system storage usage percentage. You can link to Virtual Machines workspace from Top 5 Hypervisors by Storage Usage Percentage view to check the virtual machines created by a specific hypervisor. To link to that workspace, right-click the bar that represents the hypervisor you want to check and click Link To, then click Virtual Machines.

Object Containers workspace

The Object Containers workspace is a link to workspace that displays the information about each object container, such as container name, object count and used bytes. You can link to Objects workspace from Object Containers workspace to check the object list for a specific object container. To link to that workspace, right-click the row that represents an object container and then select the link to Objects in the context menu.

This workspace contains the following view:

Object Containers

The Object Containers view shows a table view displaying the information about object containers.

Object Storage Account Statistics workspace

The Object Storage Account Statistics workspace is a link to the workspace that displays the information about the account of current project, such as container count, object count and used bytes. You can link to Object Containers workspace from Object Storage Account Statistics workspace to check the container list for the account. To link to that workspace, right-click the row that represents the account and then select the link to Object Containers in the context menu.

This workspace contains the following view:

Object Storage Account Statistics

The Object Storage Account Statistics view shows a table view displaying the information about the account of current project.

Objects workspace

The Objects workspace is a link to workspace that displays the information about each object in an object container, such as its name, size, content type and last modified time.

This workspace contains the following view:

Objects

The Objects view shows a table view displaying the information about objects.

Process Details workspace

The Process Details workspace is a link to workspace that displays the detailed information about each process, such as its process name, PID, PPID, CPU usage, service type, status.

This workspace contains the following view:

Process Details

The Process Details view shows a table view displaying the detailed information about processes.

Processes workspace

The Processes workspace displays processes status summarized by OpenStack service type or host in monitored OpenStack environment. You can link to Process Details workspace from

Processes navigation node to check the detailed information of all processes. To link to that workspace, right-click the Processes navigation node and select Process Details link.

This workspace contains the following views:

Top 5 Components by Process Count

The Top 5 Components by Process Count view displays the top 5 service types by the process count of each service type. You can link to Process Details workspace to check the detailed information of processes that belong to the selected service type. To link to that workspace, right-click the component and select Process Details link.

Top 5 Hosts by Process Count

The Top 5 Hosts by Process Count view displays the top 5 hosts by the process count of each host. You can link to Process Details workspace to check the detailed information of processes that belong to the selected host. To link to that workspace, right-click the host and select Process Details link.

Process Summary by Component

The Process Summary by Component view shows a table view displaying the number of processes of each OpenStack component. You can link to Process Details workspace to check the detailed information of processes that belong to the selected component. To link to that workspace, right-click the component and select Process Details link.

Process Summary by Host

The Process Summary by Host view shows a table view displaying the number of processes of each OpenStack host. You can link to Process Details workspace to check the detailed information of processes that belong to the selected host. To link to that workspace, right-click the host and select Process Details link.

Regions workspace

The Regions workspace displays the resource usages of the monitored OpenStack regions. You can link to Flavors and Volumes workspaces from Regions navigation node to check the flavors and volumes in all the regions. To link to that workspace, right-click the Regions navigation node and select Flavors and Volumes links.

This workspace contains the following views:

Region Summary

The Region Summary view shows a table view displaying the resource usages of all the regions, such as each region's VCPU, Memory, File System Storage usage percentage and the number of virtual machine instances. You can link to Hypervisors, Virtual Machines, Flavors and Volumes workspaces from Region Summary table view to check the hypervisors, virtual machines, flavors and volumes in a specific region. To link to those workspaces, right-click the row that represents the region you want to check then select links to Hypervisors, Virtual Machines, Flavors and Volumes in the context menu.

Top 5 Regions by VCPU Usage Percentage

The Top 5 Regions by VCPU Usage Percentage view displays the top 5 regions that have the highest VCPU usage percentage. You can link to Hypervisors, Virtual Machines, Flavors and Volumes workspaces from Top 5 Regions by VCPU Usage Percentage view to check the hypervisors, virtual machines, flavors and volumes in a specific region. To link to those workspaces, right-click the bar that represents the region you want to check and click Links To, then click Hypervisors, Virtual Machines, Flavors and Volumes.

Top 5 Regions by Memory Usage Percentage

The Top 5 Regions by Memory Usage Percentage view displays the top 5 regions that have the highest memory usage percentage. You can link to Hypervisors, Virtual Machines, Flavors and Volumes workspaces from Top 5 Regions by Memory Usage Percentage view to check the hypervisors, virtual machines, flavors and volumes in a specific region. To link to those workspaces, right-click the bar that represents the region you want to check and click Links To, then click Hypervisors, Virtual Machines, Flavors and Volumes.

Top 5 Regions by Storage Usage Percentage

The Top 5 Regions by Storage Usage Percentage view displays the top 5 regions that have the highest file system storage usage percentage. You can link to Hypervisors, Virtual

Machines, Flavors and Volumes workspaces from Top 5 Regions by Storage Usage Percentage view to check the hypervisors, virtual machines, flavors and volumes in a specific region. To link to those workspaces, right-click the bar that represents the region you want to check and click Links To, then click Hypervisors, Virtual Machines, Flavors and Volumes.

Top 5 Regions by Floating IP Quota Percentage

The Top 5 Regions by Floating IP Quota Percentage view displays the top 5 regions that have the highest floating IP quota percentage in current project. You can link to Hypervisors, Virtual Machines, Flavors and Volumes workspaces from Top 5 Regions by Floating IP Quota Percentage view to check the hypervisors, virtual machines, flavors and volumes in a specific region. To link to those workspaces, right-click the bar that represents the region you want to check and click Links To, then click Hypervisors, Virtual Machines, Flavors and Volumes.

Virtual Machines workspace

The Virtual Machines workspace is a link to workspace that displays the information about each virtual machine, such as its ID, name, IP addresses, status, image, flavor, key pairs, tenant, user and user email.

This workspace contains the following view:

Virtual Machines

The Virtual Machines view shows a table view displaying the detailed information about virtual machines. You can link to OS agent from Virtual Machines table view to check the performance of a specific virtual machine if OS agent has been installed in the virtual machine and connected to the same ITM Hub TEMS with the OpenStack monitoring agent. To link to that workspace, right-click the row that represents the virtual machine you want to check and then select the link to OS Agent in the context menu.

Volumes workspace

The Volumes workspace is a link to workspace that displays the information about each volume, such as its ID, name, type, size, metadata, status and attachments.

This workspace contains the following view:

Volumes

The Volumes view shows a table view displaying the detailed information about volumes.

Chapter 4. Attributes reference

Attributes are the application properties that are being measured and reported by the IBM Monitoring Agent for OpenStack.

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 Monitoring Agent for OpenStack 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 31.

Attribute groups for the monitoring agent

The Monitoring Agent for OpenStack 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: API Endpoint Details
 - Table name: KSGAPIDET
 - Warehouse table name: KSG_API_DETAILS or KSGAPIDET
- Attribute group name: API Endpoints
 - Table name: KSGSERSUM
 - Warehouse table name: KSG_API_STATUS_SUMMARY or KSGSERSUM
- Attribute group name: Environment Summary
 - Table name: KSGOSSUM
 - Warehouse table name: KSG_ENVIRONMENT_SUMMARY or KSGOSSUM
- Attribute group name: Flavors

- Table name: KSGFLAVORR
- Warehouse table name: KSG_FLAVORS or KSGFLAVORR
- Attribute group name: Hypervisors
 - Table name: KSGHYPERVRS
 - Warehouse table name: KSG_HYPERVISORS or KSGHYPERVRS
- Attribute group name: Object Containers
 - Table name: KSGCONTANS
 - Warehouse table name: KSG_OBJECT_CONTAINERS or KSGCONTANS
- Attribute group name: Object Storage Account Statistics
 - Table name: KSGACCOUNT
 - Warehouse table name: KSG_ACCOUNT_STATISTICS or KSGACCOUNT
- Attribute group name: Objects
 - Table name: KSGOBJECTS
 - Warehouse table name: KSG_OBJECTS
- Attribute group name: Performance Object Status
 - Table name: KSGPOBJST
 - Warehouse table name: KSG_PERFORMANCE_OBJECT_STATUS or KSGPOBJST
- Attribute group name: Process Details
 - Table name: KSGPROCDET
 - Warehouse table name: KSG_PROCESS_DETAILS or KSGPROCDET
- Attribute group name: Processes
 - Table name: KSGPROCSUM
 - Warehouse table name: KSG_PROCESS_SUMMARY or KSGPROCSUM
- Attribute group name: Processes by Host
 - Table name: KSGPRCHOST
 - Warehouse table name: KSG_PROCESS_SUMMARY_BY_HOST or KSGPRCHOST
- Attribute group name: Regions
 - Table name: KSGREGISUM
 - Warehouse table name: KSG_REGION_SUMMARY or KSGREGISUM
- Attribute group name: Take Action Status
 - Table name: KSGTACTST
 - Warehouse table name: KSG_TAKE_ACTION_STATUS or KSGTACTST
- Attribute group name: Thread Pool Status
 - Table name: KSGTHPLST
 - Warehouse table name: KSG_THREAD_POOL_STATUS or KSGTHPLST
- Attribute group name: VM Details
 - Table name: KSGINSTDET
 - Warehouse table name: KSG_VM_DETAILS or KSGINSTDET
- Attribute group name: Volumes
 - Table name: KSGVOLUMES
 - Warehouse table name: KSG_VOLUMES

Attributes in each attribute group

Attributes in each Monitoring Agent for OpenStack 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.

API Endpoint Details attribute group

The OpenStack API details

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 API Endpoint Details attribute group:

API Endpoint Detecting Failure Percent attribute

Description

The API endpoint detecting failure percentage during the last collection interval

Type

integer (32-bit gauge)

Warehouse name

API_FAILURE_PERCENT or F_PERC

API Endpoint Detecting Failure Times attribute

Description

The count of API endpoint detecting failure during the last collection interval

Type

integer (32-bit gauge)

Warehouse name

API_DETECTING_FAILURE_TIMES or F_TIMES

API Endpoint Detecting Status Code attribute

Description

The return code of API endpoint detecting

Type

integer (32-bit gauge)

Warehouse name

API_STATUS_CODE or SCODE

API Endpoint Detecting Times attribute

Description

The count of API endpoint detecting during the last collection interval

Type

integer (32-bit gauge)

	Warehouse name API_DETECTING_TIMES or D_TIMES
<u>Keystone Host attribute</u>	This attribute is a key attribute.
	Description The host of Keystone service
	Type string
	Warehouse name KEYSTONE_HOST or KS_HOST
<u>Node attribute</u>	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
<u>Public URL attribute</u>	This attribute is a key attribute.
	Description The public URL of the API endpoint
	Type string
	Warehouse name PUBLIC_URL or API_URL
<u>Region Name attribute</u>	This attribute is a key attribute.
	Description The name of the monitored region
	Type string
	Warehouse name REGION_NAME or REGI_NAME
<u>Service Type attribute</u>	
	Description The service type of the API endpoint
	Type string
	Warehouse name COMPONENT_NAME or COMP_NAME
<u>Status attribute</u>	
	Description The current status of the API endpoint
	Type integer with enumerated values. The following values are defined: Available (0), Unavailable (1), Unreachable (2). Any value that does not have a definition here is displayed in the User Interface
	Warehouse name STATUS
<u>Timestamp attribute</u>	
	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

API Endpoints attribute group

The OpenStack API status summary by component

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 API Endpoints attribute group:

Available API Endpoints attribute

Description

The number of available API endpoints

Type

integer (32-bit gauge)

Warehouse name

AVAILABLE_API_ENDPOINTS or AVAI_API

Keystone Host attribute This attribute is a key attribute.

Description

The host of Keystone service

Type

string

Warehouse name

KEYSTONE_HOST or KS_HOST

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

Service Type attribute This attribute is a key attribute.

Description

The service type of the API endpoint

Type

string

Warehouse name

COMPONENT_NAME or COMP_NAME

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

Unavailable API Endpoints attribute

Description

The number of unavailable API endpoints

Type

integer (32-bit gauge)

Warehouse name

UNAVAILABLE_API_ENDPOINTS or UNAVA_API

Environment Summary attribute group

The summary information of the OpenStack environment

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 Environment Summary attribute group:

Active Instances attribute

Description

The number of power-on OpenStack virtual machine instances

Type

integer (32-bit gauge)

Warehouse name

ACTIVE_INSTANCE_COUNT or ACT_INST

Available API Endpoints attribute

Description

The number of available OpenStack API Endpoints

Type

integer (32-bit gauge)

Warehouse name

AVAILABLE_API_COUNT or AVAI_API

Inactive Instances attribute

Description

The number of power-off OpenStack virtual machine instances

Type

integer (32-bit gauge)

Warehouse name

INACTIVE_INSTANCE_COUNT or INAC_INST

Keystone Host attribute This attribute is a key attribute.

Description

The host of Keystone service

Type

string

Warehouse name

KEYSTONE_HOST or KS_HOST

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

Region Count (Critical File System Storage Usage) attribute

Description

The number of regions with critical file system storage usage (Greater than or equal 90%)

Type

integer (32-bit gauge)

Warehouse name

FILE_SYSTEM_STORAGE_USAGE_CRITICAL_REGIONS or FSS_RED

Region Count (Critical Floating IP Quota Usage) attribute

Description

The number of regions with critical floating IP quota usage for current project (Greater than or equal 90%)

Type

integer (32-bit gauge)

Warehouse name

FLOATING_IP_USAGE_CRITICAL_REGIONS or IP_RED

Region Count (Critical Memory Usage) attribute**Description**

The number of regions with critical memory usage (Greater than or equal 90%)

Type

integer (32-bit gauge)

Warehouse name

MEMORY_USAGE_CRITICAL_REGIONS or MEM_RED

Region Count (Critical vCPU Usage) attribute**Description**

The number of regions with critical vCPU usage (Greater than or equal 90%)

Type

integer (32-bit gauge)

Warehouse name

VCPU_USAGE_CRITICAL_REGIONS or VCPU_RED

Region Count (Normal File System Storage Usage) attribute**Description**

The number of regions with normal file system storage usage (Less than 80%)

Type

integer (32-bit gauge)

Warehouse name

FILE_SYSTEM_STORAGE_USAGE_NORMAL_REGIONS or FSS_GRE

Region Count (Normal Floating IP Quota Usage) attribute**Description**

The number of regions with normal floating IP quota usage for current project (Less than 80%)

Type

integer (32-bit gauge)

Warehouse name

FLOATING_IP_USAGE_NORMAL_REGIONS or IP_GRE

Region Count (Normal Memory Usage) attribute**Description**

The number of regions with normal memory usage (Less than 80%)

Type

integer (32-bit gauge)

Warehouse name

MEMORY_USAGE_NORMAL_REGIONS or MEM_GRE

Region Count (Normal vCPU Usage) attribute**Description**

The number of regions with normal vCPU usage (Less than 80%)

Type

integer (32-bit gauge)

Warehouse name

VCPU_USAGE_NORMAL_REGIONS or VCPU_GRE

Region Count (Warning File System Storage Usage) attribute**Description**

The number of regions with warning file system storage usage (Greater than or equal 80% and less than 90%)

Type

integer (32-bit gauge)

Warehouse name

FILE_SYSTEM_STORAGE_USAGE_WARNING_REGIONS or FSS_YEL

Region Count (Warning Floating IP Quota Usage) attribute**Description**

The number of regions with warning floating IP quota usage for current project (Greater than or equal 80% and less than 90%)

Type

integer (32-bit gauge)

Warehouse name

FLOATING_IP_USAGE_WARNING_REGIONS or IP_YEL

Region Count (Warning Memory Usage) attribute**Description**

The number of regions with warning memory usage (Greater than or equal 80% and less than 90%)

Type

integer (32-bit gauge)

Warehouse name

MEMORY_USAGE_WARNING_REGIONS or MEM_YEL

Region Count (Warning vCPU Usage) attribute**Description**

The number of regions with warning vCPU usage (Greater than or equal 80% and less than 90%)

Type

integer (32-bit gauge)

Warehouse name

VCPU_USAGE_WARNING_REGIONS or VCPU_YEL

Running Processes attribute**Description**

The number of running OpenStack processes

Type

integer (32-bit gauge)

Warehouse name

RUNNING_PROCESS_COUNT or AVAIL_PROC

Tenant attribute This attribute is a key attribute.**Description**

The tenant to which the current user belongs

Type

string

Warehouse name

TENANT_NAME or TENNAME

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

Unavailable API Endpoints attribute**Description**

The number of unavailable OpenStack API Endpoints

Type

integer (32-bit gauge)

Warehouse name

UNAVAILABLE_API_COUNT or UNAV_API

Flavors attribute group

The flavors information

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 Flavors attribute group:

Ephemeral Disk (GB) attribute

Description

The ephemeral disk setting of the flavor

Type

integer (32-bit gauge)

Warehouse name

EPHEMERAL_DISK or EPHE_DISK

Flavor ID attribute This attribute is a key attribute.

Description

The flavor ID

Type

string

Warehouse name

FLAVOR_ID

Flavor Name attribute This attribute is a key attribute.

Description

The flavor name

Type

string

Warehouse name

FLAVOR_NAME or FLAV_NAME

Is Public attribute

Description

This attribute shows whether the flavor is public or not

Type

string

Warehouse name

IS_PUBLIC

Keystone Host attribute This attribute is a key attribute.

Description

The host of Keystone service of the OpenStack environment

Type

string

Warehouse name

KEYSTONE_HOST or KS_HOST

Memory (MB) attribute

Description

The memory setting of the flavor

Type

integer (32-bit gauge)

Warehouse name

MEMORY

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

Reference Count attribute**Description**

The reference count. This attribute indicates how many VMs were created with this flavor

Type

integer (32-bit gauge)

Warehouse name

REFERENCE_COUNT or REF_COUNT

Region Name attribute This attribute is a key attribute.**Description**

The name of the monitored region

Type

string

Warehouse name

REGION_NAME or REGI_NAME

Root Disk (GB) attribute**Description**

The root disk setting of the flavor

Type

integer (32-bit gauge)

Warehouse name

ROOT_DISK

Swap Disk (MB) attribute**Description**

The swap disk setting of the flavor

Type

integer (32-bit gauge)

Warehouse name

SWAP_DISK

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

vCPU Count attribute**Description**

The vCPU setting of the flavor

Type

integer (32-bit gauge)

Warehouse name

VCPU_COUNT or VCPU

Hypervisors attribute group

The hypervisor list

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 Hypervisors attribute group:

CPU Architecture attribute

Description
The hypervisor CPU architecture

Type
string

Warehouse name
CPU_ARCHITECTURE or CPU_ARCH

CPU Model attribute

Description
The hypervisor CPU model

Type
string

Warehouse name
CPU_MODEL or CPU_MOD

CPU Vendor attribute

Description
The hypervisor CPU vendor

Type
string

Warehouse name
CPU_VENDOR or CPU_VEND

Host IP attribute This attribute is a key attribute.

Description
The IP address of the host

Type
string

Warehouse name
HOST_IP

Hypervisor Hostname attribute This attribute is a key attribute.

Description
The hypervisor host name

Type
string

Warehouse name
HYPERVISOR_HOSTNAME or HPVHOST

Hypervisor ID attribute This attribute is a key attribute.

Description
The hypervisor ID

Type
string

Warehouse name
ID or HYPERV_ID

Hypervisor Type attribute

Description
The hypervisor type

Type
string

Warehouse name
HYPERVISOR_TYPE or HPVTYPE

Hypervisor Version attribute

Description
The hypervisor version

Type
string

Warehouse name
HYPERVISOR_VERSION or HPVVER

Keystone Host attribute This attribute is a key attribute.

Description
The host of Keystone service of the OpenStack environment

Type
string

Warehouse name
KEYSTONE_HOST or KS_HOST

Local Storage Size (TB) attribute

Description
The local file system storage size (TB) of the hypervisor

Type
string

Warehouse name
LOCAL_STORAGE_SIZE or STORAGE

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

Region Name attribute This attribute is a key attribute.

Description
The name of the monitored region

Type
string

Warehouse name
REGION_NAME or REGI_NAME

Running VMs attribute

Description
The count of running virtual machine instances of the hypervisor

Type
integer (32-bit gauge)

Warehouse name
RUNNING_VMS or VMS

Service Host attribute

Description
The compute service host

Type
string

Warehouse name
SERVICE_HOST or SRVHOST

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

Total Memory Size (GB) attribute

Description
The total memory size (GB) of the hypervisor

Type	string
Warehouse name	TOTAL_MEMORY_SIZE or MEM
<u>Total vCPU Count attribute</u>	
Description	The count of all vCPUs of the hypervisor
Type	integer (32-bit gauge)
Warehouse name	TOTAL_VCPU_COUNT or VCPU
<u>Used Memory Percent attribute</u>	
Description	The percentage of memory usage of the hypervisor
Type	integer (32-bit gauge)
Warehouse name	USED_MEMORY_PERCENT or MEM_RATE
<u>Used Memory Size (GB) attribute</u>	
Description	The used memory size (GB) of the hypervisor
Type	string
Warehouse name	USED_MEMORY_SIZE or USED_MEM
<u>Used Storage Percent attribute</u>	
Description	The percentage of used local file system storage of the hypervisor
Type	integer (32-bit gauge)
Warehouse name	USED_STORAGE_PERCENT or FSTO_RATE
<u>Used Storage Size (TB) attribute</u>	
Description	The used local file system storage size (TB) of the hypervisor
Type	string
Warehouse name	USED_STORAGE_SIZE or USED_FSTO
<u>Used vCPU Count attribute</u>	
Description	The count of used vCPUs of the hypervisor
Type	integer (32-bit gauge)
Warehouse name	USED_VCPU_COUNT or USED_VCPU
<u>Used vCPU Percent attribute</u>	
Description	The percentage of vCPU usage of the hypervisor
Type	integer (32-bit gauge)
Warehouse name	USED_VCPU_PERCENT or VCPU_RATE

Object Containers attribute group

The object containers

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 Object Containers attribute group:

Accept Ranges attribute

Description

The accept ranges of the object storage account

Type

string

Warehouse name

ACCEPT_RANGES or ACC_RAG

Container Name attribute This attribute is a key attribute.

Description

The name of the object container

Type

string

Warehouse name

CONTAINER_NAME or CON_NAME

Content Type attribute

Description

The content type of object storage

Type

string

Warehouse name

CONTENT_TYPE or CON_TYPE

Keystone Host attribute This attribute is a key attribute.

Description

The host of Keystone service

Type

string

Warehouse name

KEYSTONE_HOST or KS_HOST

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

Object Count attribute

Description

The number of objects in the container

Type

integer (32-bit gauge)

Warehouse name

OBJECT_COUNT or OBJ_COUNT

Tenant attribute This attribute is a key attribute.

Description

The tenant to which the object storage account belongs

Type
string
Warehouse name
TENANT_NAME or TENNAME

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

Used Bytes attribute

Description
The bytes used by the container
Type
integer (64-bit gauge)
Warehouse name
USED_BYTES or USED_BYTE

Object Storage Account Statistics attribute group

The object storage account statistics

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 Object Storage Account Statistics attribute group:

Accept Ranges attribute

Description
The accept ranges of the object storage account
Type
string
Warehouse name
ACCEPT_RANGES or ACC_RAG

Container Count attribute

Description
The number of containers
Type
integer (32-bit gauge)
Warehouse name
CONTAINER_COUNT or CON_COUNT

Content Type attribute

Description
The content type of object storage
Type
string
Warehouse name
CONTENT_TYPE or CON_TYPE

Keystone Host attribute This attribute is a key attribute.

Description
The host of Keystone service
Type
string

Warehouse name	KEYSTONE_HOST or KS_HOST
<u>Node attribute</u>	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
<u>Object Count attribute</u>	
Description	The number of objects
Type	integer (32-bit gauge)
Warehouse name	OBJECT_COUNT or OBJ_COUNT
<u>Tenant attribute</u>	This attribute is a key attribute.
Description	The tenant to which the object storage account belongs
Type	string
Warehouse name	TENANT_NAME or TENNAME
<u>Timestamp attribute</u>	
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
<u>Used Bytes attribute</u>	
Description	The bytes used by the account
Type	integer (64-bit gauge)
Warehouse name	USED_BYTES or USED_BYTE

Objects attribute group

The objects in containers

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 Objects attribute group:

Bytes attribute

Description	The length of the object in bytes
Type	integer (64-bit gauge)
Warehouse name	BYTES or LENGTH

Container Name attribute This attribute is a key attribute.

Description

The name of the object container

Type

string

Warehouse name

CONTAINER_NAME or CON_NAME

Content Type attribute

Description

The content type of object storage

Type

string

Warehouse name

CONTENT_TYPE or CON_TYPE

Keystone Host attribute This attribute is a key attribute.

Description

The host of Keystone service

Type

string

Warehouse name

KEYSTONE_HOST or KS_HOST

Last Modified attribute

Description

The last modified time of the object

Type

string

Warehouse name

LAST_MODIFIED or MOD_TIME

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

Object Name attribute This attribute is a key attribute.

Description

The name of the object

Type

string

Warehouse name

OBJECT_NAME or OBJECT

Tenant attribute This attribute is a key attribute.

Description

The tenant to which the object storage account belongs

Type

string

Warehouse name

TENANT_NAME or TENNAME

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

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 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, so you can see whether the agent is collecting data 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:

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 following values are defined: NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

AVERAGE_COLLECTION_DURATION or COLAVGD

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

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

Error Code attribute**Description**

The error code associated with the query.

Type

integer with enumerated values. 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 value that does not have a definition here is displayed in the User Interface

Warehouse name

ERROR_CODE or ERRCODE

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

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

Last Collection Finished attribute**Description**

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

Type
timestamp with enumerated values. The following values are defined: NOT COLLECTED (0691231190000000), NOT COLLECTED (0000000000000001). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
LAST_COLLECTION_FINISHED or COLFINI

Last Collection Start attribute

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

Type
timestamp with enumerated values. The following values are defined: NOT COLLECTED (0691231190000000), NOT COLLECTED (0000000000000001). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
LAST_COLLECTION_START or COLSTRT

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

Number of Collections attribute

Description
The number of data collections for this group since the agent started.

Type
integer (32-bit counter)

Warehouse name
NUMBER_OF_COLLECTIONS or NUMCOLL

Object Name attribute

Description
The name of the performance object.

Type
string

Warehouse name
OBJECT_NAME or OBJNAME

Object Status attribute

Description
The status of the performance object.

Type
integer with enumerated values. The following values are defined: ACTIVE (0), INACTIVE (1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
OBJECT_STATUS or OBJSTTS

Object Type attribute

Description
The type of the performance object.

Type
integer with enumerated values. 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 value that does not have a definition here is displayed in the User Interface

Warehouse name

OBJECT_TYPE or OBJTYPE

Query Name attribute This attribute is a key attribute.

Description

The name of the attribute group.

Type

string

Warehouse name

QUERY_NAME or ATTRGRP

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

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

Process Details attribute group

The OpenStack process details

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 Process Details attribute group:

Component Name attribute

Description

The OpenStack component name

Type

string

Source

The source for this attribute is Script data.

Warehouse name

COMPONENT_NAME or COMP_NAME

CPU Usage(%) attribute This attribute is a key attribute.

Description

The percentage of CPU usage of a specific OpenStack process

Type

real number (32-bit gauge) with one decimal places of precision

Source

The source for this attribute is Script data.

Warehouse name

CPU

Host IP attribute This attribute is a key attribute.

Description

The IP of the host where the process is running on

Type

string

Source

The source for this attribute is Script data.

Warehouse name

HOST_IP

Host Name attribute This attribute is a key attribute.**Description**

The host where the process is running on

Type

string

Source

The source for this attribute is Script data.

Warehouse name

HOSTNAME

Keystone Host attribute This attribute is a key attribute.**Description**

The host of Keystone service

Type

string

Warehouse name

KEYSTONE_HOST or KS_HOST

Memory Usage(%) attribute This attribute is a key attribute.**Description**

The percentage of memory usage of a specific OpenStack process

Type

real number (32-bit gauge) with one decimal places of precision

Source

The source for this attribute is Script data.

Warehouse name

MEM

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

pid attribute This attribute is a key attribute.**Description**

The process ID

Type

integer (32-bit numeric property)

Source

The source for this attribute is Script data.

Warehouse name

PID

ppid attribute This attribute is a key attribute.**Description**

The parent process ID

Type

integer (32-bit numeric property)

Source
The source for this attribute is Script data.

Warehouse name
PPID

Process Name attribute This attribute is a key attribute.

Description
The process name

Type
string

Source
The source for this attribute is Script data.

Warehouse name
PROCESS_NAME or PNAME

Start Time attribute

Description
The specific OpenStack process start time.

Type
string

Source
The source for this attribute is Script data.

Warehouse name
START_TIME

Status attribute

Description
The current status of a specific OpenStack process

Type
integer with enumerated values. The following values are defined: Running (0), Defunct (1). Any value that does not have a definition here is displayed in the User Interface

Source
The source for this attribute is Script data.

Warehouse name
STATUS

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

Processes attribute group

The OpenStack process summary by service

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 Processes attribute group:

Component Name attribute This attribute is a key attribute.

Description
The component name

Type
string

Source
The source for this attribute is Script data.

Warehouse name
COMPONENT_NAME or COMP_NAME

Component Process Count attribute This attribute is a key attribute.

Description
The component process count

Type
integer (32-bit gauge)

Source
The source for this attribute is Script data.

Warehouse name
COMPONENT_PROCESS_COUNT or COMPPRCNUM

Keystone Host attribute This attribute is a key attribute.

Description
The host of Keystone service

Type
string

Warehouse name
KEYSTONE_HOST or KS_HOST

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

Processes by Host attribute group

The OpenStack process summary by host

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 Processes by Host attribute group:

Host Name attribute This attribute is a key attribute.

Description
The host name

Type
string

Source
The source for this attribute is Script data.

Warehouse name
HOSTNAME

Host Process Count attribute This attribute is a key attribute.

Description

The host process count

Type

integer (32-bit gauge)

Source

The source for this attribute is Script data.

Warehouse name

HOST_PROCESS_COUNT or HOSTPRCNUM

Keystone Host attribute This attribute is a key attribute.

Description

The host of Keystone service

Type

string

Warehouse name

KEYSTONE_HOST or KS_HOST

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

SSH Connection Status

Description

The status of the ssh connection to target host.

Type

string

Warehouse name

SSH_Status or SSHSTATUS

Regions attribute group

The OpenStack region resource usage summary

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 Regions attribute group:

Active Instances attribute

Description

The number of power-on instances in monitored OpenStack region

Type

integer (32-bit gauge)

Warehouse name

ACTIVE_INSTANCE_COUNT or ACT_INST

Allocated Floating IP Count attribute

Description

The number of allocated floating IP addresses for current project

Type

integer (32-bit gauge)

Warehouse name

ALLOCATED_FLOATING_IP_COUNT or USED_IP

Floating IP Quota attribute

Description

The floating IP quota for current project

Type

integer (32-bit gauge)

Warehouse name

FLOATING_IP_QUOTA or IP_QUOTA

Inactive Instances attribute

Description

The number of power-off instances in monitored OpenStack region

Type

integer (32-bit gauge)

Warehouse name

INACTIVE_INSTANCE_COUNT or INAC_INST

Keystone Host attribute This attribute is a key attribute.

Description

The host of Keystone service

Type

string

Warehouse name

KEYSTONE_HOST or KS_HOST

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

Region Name attribute This attribute is a key attribute.

Description

The name of the monitored region

Type

string

Warehouse name

REGION_NAME or REGI_NAME

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

Total File System Storage Size (TB) attribute

Description

The total file system storage Size (TB) of the monitored OpenStack region

Type	string
Warehouse name	TOTAL_FILE_SYSTEM_STORAGE_SIZE or STORAGE
<u>Total Floating IP Count attribute</u>	
Description	The number of floating IP addresses for current project
Type	integer (32-bit gauge)
Warehouse name	TOTAL_FLOATING_IP_COUNT or FLOATIP
<u>Total Instances attribute</u>	
Description	The number of total instances in monitored OpenStack region
Type	integer (32-bit gauge)
Warehouse name	TOTAL_INSTANCE_COUNT or INST
<u>Total Memory Size (GB) attribute</u>	
Description	The total memory size (GB) in the monitored OpenStack region
Type	string
Warehouse name	TOTAL_MEMORY_SIZE or MEM
<u>Total vCPUs attribute</u>	
Description	The number of total vCPUs in the monitored OpenStack region
Type	integer (32-bit gauge)
Warehouse name	TOTAL_VCPU_COUNT or VCPU
<u>Used File System Storage Percent attribute</u>	
Description	The percentage of used file system storage in the monitored OpenStack region
Type	integer (32-bit gauge)
Warehouse name	USED_FILE_SYSTEM_STORAGE_PERCENT or FSTO_RATE
<u>Used File System Storage Size (TB) attribute</u>	
Description	The used file system storage Size (TB) of the monitored OpenStack region
Type	string
Warehouse name	USED_FILE_SYSTEM_STORAGE_SIZE or USED_FSTO
<u>Used Floating IP Quota Percent attribute</u>	
Description	The percentage of used floating IP quota for the current project
Type	integer (32-bit gauge)
Warehouse name	ALLOCATED_FLOATING_IP_PERCENT or UIP_RATE
<u>Used Memory Percent attribute</u>	
Description	The percentage of memory usage of the monitored OpenStack region

Type
integer (32-bit gauge)

Warehouse name
USED_MEMORY_PERCENT or MEM_RATE

Used Memory Size (GB) attribute

Description
The used memory size (GB) in the monitored OpenStack region

Type
string

Warehouse name
USED_MEMORY_SIZE or USED_MEM

Used vCPU Percent attribute

Description
The percentage of vCPU usage of the monitored OpenStack region

Type
integer (32-bit gauge)

Warehouse name
USED_VCPU_PERCENT or VCPU_RATE

Used vCPUs attribute

Description
The number of used vCPUs in the monitored OpenStack region

Type
integer (32-bit gauge)

Warehouse name
USED_VCPU_COUNT or USED_VCPU

Take Action Status attribute group

The Take Action Status data set contains information about the results of actions this agent ran.

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 Take Action Status attribute group:

Action App Return Code attribute

Description
The application return code of the action is the actual return code of the command run.

Type
integer (32-bit numeric property)

Warehouse name
ACTION_APP_RETURN_CODE or TSKAPRC

Action Command attribute

Description
The command that was run by the action.

Type
string

Warehouse name
ACTION_COMMAND or TSKCMND

Action ID attribute

Description
The id of the action.

Type
integer (32-bit numeric property)

Warehouse name
ACTION_ID or TSKID

Action Instance attribute

Description

The instance that is associated with the output produced by running the action.

Type

string

Warehouse name

ACTION_INSTANCE or TSKINST

Action Message attribute

Description

The message that is associated with the return code of the action.

Type

string

Warehouse name

ACTION_MESSAGE or TSKMSGE

Action Name attribute

Description

The name of the action that was run.

Type

string

Warehouse name

ACTION_NAME or TSKNAME

Action Node attribute

Description

The node where the action ran.

Type

string

Warehouse name

ACTION_NODE or TSKORGN

Action Owner attribute

Description

The name of the situation or user that initiated the action.

Type

string

Warehouse name

ACTION_OWNER or TSKOWNR

Action Results attribute

Description

The output that is produced by running the action.

Type

string

Warehouse name

ACTION_RESULTS or TSKOUTP

Action Status attribute

Description

The return code from the Action Status dialog, which is the return code category that is defined for the application return code.

Type

integer (32-bit numeric property) with enumerated values. The following values are defined: OK (0), NOT APPLICABLE (1), GENERAL ERROR (2), WARNING (3), NOT RUNNING (4), DEPENDENT NOT RUNNING (5), ALREADY RUNNING (6), PREREQ NOT RUNNING (7), TIMED OUT (8), DOESNT EXIST (9), UNKNOWN (10), DEPENDENT STILL RUNNING (11), INSUFFICIENT USER AUTHORITY (12). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

ACTION_STATUS or TSKSTAT

Action Subnode attribute

Description

The subnode where the action ran.

Type

string

Warehouse name

ACTION_SUBNODE or TSKSBND

Action Type attribute

Description

The type of the action.

Type

integer (32-bit numeric property) with enumerated values. The following values are defined: UNKNOWN (0), AUTOMATION (1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

ACTION_TYPE or TSKTYPE

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 time the event was generated.

Type

string

Source

The source for this attribute is the agent.

Warehouse name

TIMESTAMP

Thread Pool Status attribute group

The Thread Pool Status attribute group contains information that reflects the status of the internal thread pool used to collect data asynchronously.

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 Thread Pool 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

Thread Pool Active Threads attribute

Description

The number of threads in the thread pool currently active doing work.

Type
integer (32-bit gauge) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_ACTIVE_THREADS or TPACTTH

Thread Pool Avg Active Threads attribute

Description
The average number of threads in the thread pool simultaneously active doing work.

Type
real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_AVG_ACTIVE_THREADS or TPAVGAT

Thread Pool Avg Job Wait attribute

Description
The average time a job spends waiting on the thread pool queue in seconds.

Type
real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_AVG_JOB_WAIT or TPAVJBW

Thread Pool Avg Queue Length attribute

Description
The average length of the thread pool queue during this run.

Type
real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_AVG_QUEUE_LENGTH or TPAVGQL

Thread Pool Max Active Threads attribute

Description
The peak number of threads in the thread pool that were simultaneously active doing work.

Type
integer (32-bit counter) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_MAX_ACTIVE_THREADS or TPMAXAT

Thread Pool Max Queue Length attribute

Description
The peak length the thread pool queue reached.

Type
integer (32-bit counter) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_MAX_QUEUE_LENGTH or TPMAXQL

Thread Pool Max Size attribute

Description
The maximum number of threads that are allowed to exist in the thread pool.

Type
integer (32-bit numeric property) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_MAX_SIZE or TPMAXSZ

Thread Pool Min Active Threads attribute

Description
The smallest number of threads in the thread pool that were simultaneously active doing work.

Type
integer (32-bit counter) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_MIN_ACTIVE_THREADS or TPMINAT

Thread Pool Min Queue Length attribute

Description
The minimum length the thread pool queue reached.

Type
integer (32-bit counter) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_MIN_QUEUE_LENGTH or TPMINQL

Thread Pool Queue Length attribute

Description
The number of jobs currently waiting in the thread pool queue.

Type
integer (32-bit gauge) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_QUEUE_LENGTH or TPQLGTH

Thread Pool Size attribute

Description
The number of threads currently existing in the thread pool.

Type
integer (32-bit gauge) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_SIZE or THPSIZE

Thread Pool Total Jobs attribute

Description
The number of jobs that are completed by all threads in the pool since agent start.

Type
integer (32-bit counter) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
THREAD_POOL_TOTAL_JOBS or TPTJOBS

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

VM Details attribute group

The virtual machine instances

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 VM Details attribute group:

Availability Zone attribute

Description The availability zone of the virtual machine instance
Type string
Warehouse name AVAILABILITY_ZONE or ZONE

Created Time attribute

Description The virtual machine instance created time
Type string
Warehouse name CREATED_TIME or CRT_TIME

Flavor attribute

Description The flavor that is referenced by virtual machine instance
Type string
Warehouse name FLAVOR

Host attribute

Description The host where the virtual machine instance is running on
Type string
Warehouse name HOST

Hypervisor Hostname attribute

Description The hypervisor host name
Type string
Warehouse name HYPERVISOR_HOSTNAME or HPVHOST

Image attribute

Description The image used by virtual machine instance
Type string
Warehouse name IMAGE

Instance ID attribute This attribute is a key attribute.

Description

The virtual machine instance ID

Type

string

Warehouse name

INSTANCE_ID or INS_ID

Instance Name attribute

Description

The virtual machine instance name

Type

string

Warehouse name

INSTANCE_NAME or INS_NAME

IP Addresses attribute

Description

The virtual machine instance IP addresses

Type

string

Warehouse name

IP_ADDRESSES or IP_ADDR

IP Key attribute

Description

The selected instance ip addresses

Type

string

Warehouse name

IP_KEY or IPKEY

Key Name attribute

Description

The key pair name which is used to access the virtual machine

Type

string

Warehouse name

KEY_NAME or KEYNAME

Keystone Host attribute This attribute is a key attribute.

Description

The host of Keystone service

Type

string

Warehouse name

KEYSTONE_HOST or KS_HOST

Launched Time attribute

Description

The virtual machine instance launched time

Type

string

Warehouse name

LAUNCHED_TIME or LAU_TIME

Name attribute

Description

The virtual machine instance display name

Type

string

Warehouse name

NAME

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

Region Name attribute This attribute is a key attribute.

Description

The name of the monitored region

Type

string

Warehouse name

REGION_NAME or REGI_NAME

Status attribute

Description

The virtual machine instance status

Type

string

Warehouse name

STATUS or VM_STATUS

Tenant attribute

Description

The tenant to which the virtual machine instance belongs

Type

string

Warehouse name

TENANT_NAME or TENNAME

Tenant ID attribute

Description

The tenant ID which the virtual machine instance belongs to

Type

string

Warehouse name

TENANT_ID or TENID

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

Update Time attribute

Description

The virtual machine updated time

Type

string

Warehouse name

UPDATE_TIME or UPD_TIME

User attribute

Description

The user who created this virtual machine instance

Type
string
Warehouse name
USER_NAME or USERNAME

User Email attribute

Description
The email address of the virtual machine creator

Type
string

Warehouse name
USER_EMAIL or EMAIL

User ID attribute

Description
The user ID who created this virtual machine instance

Type
string

Warehouse name
USER_ID

Volumes attribute group

The volumes information

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 Volumes attribute group:

Attachments attribute

Description
The volume attachments information

Type
string

Warehouse name
ATTACHMENTS or ATTACHS

Availability Zone attribute

Description
The availability zone of the volume

Type
string

Warehouse name
AVAILABILITY_ZONE or ZONE

Bootable attribute

Description
This attribute shows whether the volume is bootable or not

Type
string

Warehouse name
BOOTABLE

Encrypted attribute

Description
This attribute shows whether the volume is encrypted or not

Type
string

Warehouse name
ENCRYPTED

Host attribute

Description

The host on which the volume is created

Type

string

Warehouse name

HOST

Keystone Host attribute This attribute is a key attribute.**Description**

The host of Keystone service of the OpenStack environment

Type

string

Warehouse name

KEYSTONE_HOST or KS_HOST

Metadata attribute**Description**

The volume metadata

Type

string

Warehouse name

METADATA

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

Project attribute**Description**

The project to which the volume belongs

Type

string

Warehouse name

PROJECT

Project ID attribute This attribute is a key attribute.**Description**

The project ID which the volume belongs to

Type

string

Warehouse name

PROJECT_ID or PROID

Region Name attribute This attribute is a key attribute.**Description**

The name of the monitored region

Type

string

Warehouse name

REGION_NAME or REGI_NAME

Size (GB) attribute**Description**

The volume size (GB)

Type

integer (32-bit gauge)

	Warehouse name
	SIZE
Status attribute	
	Description
	The volume status
	Type
	string
	Warehouse name
	STATUS
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
Type attribute	
	Description
	The volume type
	Type
	string
	Warehouse name
	TYPE
Volume ID attribute	This attribute is a key attribute.
	Description
	The volume ID
	Type
	string
	Warehouse name
	VOLUME_ID
Volume Name attribute	
	Description
	The volume display name
	Type
	string
	Warehouse name
	VOLUME_NAME or VOL_NAME

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 29.

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 29.

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 2. Capacity planning for historical data logged by the Monitoring Agent for OpenStack

Table	Attribute group	Bytes per row (agent)	Database bytes per row (warehouse)	Aggregate bytes per row (warehouse)
KSGACCOUNT	KSG_ACCOUNT_STATISTICS	348	351	509
KSGAPIDET	KSG_API_DETAILS	416	421	614
KSGSERSUM	KSG_API_STATUS_SUMMARY	180	180	295
KSGOSSUM	KSG_ENVIRONMENT_SUMMARY	272	287	987
KSGFLAVORR	KSG_FLAVORS	420	427	698
KSGHYPERVS	KSG_HYPERVISORS	932	949	1220
KSGOBJECTS	KSG_OBJECTS	532	535	615
KSGCONTANS	KSG_OBJECT_CONTAINERS	408	411	530
KSGPOBJST	KSG_PERFORMANCE_OBJECT_STATUS	352	399	664
KSGPROCDET	KSG_PROCESS_DETAILS	480	511	548
KSGPROCSUM	KSG_PROCESS_SUMMARY	208	207	244
KSGPRCHOST	KSG_PROCESS_SUMMARY_BY_HOST	208	207	244
KSGREGISUM	KSG_REGION_SUMMARY	380	394	899
KSGTACTST	KSG_TAKE_ACTION_STATUS	3480	3512	3549
KSGTHPLST	KSG_THREAD_POOL_STATUS	124	168	550
KSGINSTDET	KSG_VM_DETAILS	1692	1710	1747
KSGVOLUMES	KSG_VOLUMES	928	939	1015

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 Monitoring Agent for OpenStack. 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 Monitoring Agent for OpenStack 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.

APIEndpoint_Unavailable

- The API Endpoint is unavailable.
- Severity: Warning

- Sampling interval: 300 seconds

Keystone_Unreachable

- Unable to connect to Keystone service.
- Severity: Critical
- Sampling interval: 180 seconds

Hyperv_vCPU_High_Warn

- The used vCPU percentage of the hypervisor has exceeded 1000%.
- Severity: Warning
- Sampling interval: 315 seconds

Hyperv_vCPU_High_Cri

- The used vCPU percentage of the hypervisor has exceeded 1500%.
- Severity: Critical
- Sampling interval: 300 seconds

Hyperv_Memory_High_Warn

- The used memory percentage of the hypervisor has exceeded 100%.
- Severity: Warning
- Sampling interval: 320 seconds

Hyperv_Memory_High_Cri

- The used memory percentage of the hypervisor has exceeded 150%.
- Severity: Critical
- Sampling interval: 305 seconds

Hyperv_Storage_High_Warn

- The used local file system storage percentage of the hypervisor has exceeded 80%.
- Severity: Warning
- Sampling interval: 325 seconds

Hyperv_Storage_High_Cri

- The used local file system storage percentage of the hypervisor has exceeded 90%.
- Severity: Critical
- Sampling interval: 310 seconds

SSHConnection_Failed

- Failed to connect to target host to collect OpenStack process information.
- Severity: Warning
- Sampling interval: 300 seconds

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.

OpenStack Navigator item

No predefined situations are included for this Navigator item.

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 Monitoring Agent for OpenStack 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 Monitoring Agent for OpenStack 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 102.

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 “Prerequisites” on page 5 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 78
- “Consulting the lists of identified problems and workarounds” on page 78

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 3.

Table 3. 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 79 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> .
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 Monitoring Agent for OpenStack
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 “`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 Monitoring Agent for OpenStack. *Logging* refers to the text messages and trace data that is generated by the Monitoring Agent for OpenStack. 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 88.

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 79
- “Examples: Using trace logs” on page 82
- “Setting RAS trace parameters by using the GUI” on page 83

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 4 provides the names, locations, and descriptions of IBM Tivoli Monitoring general RAS1 log files. The log file names for the Monitoring Agent for OpenStack 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 Monitoring Agent for OpenStack, the product code is sg.

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

Table 4. 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 IBM Tivoli Monitoring <i>timestamp.log</i> file 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 4. 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 teps_odbc.log 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 4. Trace log files for troubleshooting agents (continued)

System where log is located	File name and path	Description
<p>On the computer that hosts the monitoring agent</p>	<p>The RAS1 log files are as follows:</p> <ul style="list-style-type: none"> • UNIX: <i>hostname_sg_instance_name_ksgagent_HEXtimestamp-nn.log</i> in the <i>install_dir/logs</i> directory • Linux: <i>hostname_sg_instance_name_ksgagent_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_sg_timestamp.log</i> - <i>hostname_sg_timestamp.pidnnnn</i> in the <i>install_dir/logs</i> path, where <i>nnnnn</i> is the process ID number 	<p>Traces activity of the monitoring agent.</p>
<p>On the computer that hosts the monitoring agent</p>	<p>The agent operations log files are as follows:</p> <p><i>instance_hostnameSG.LG0</i> is the current log created when the agent was started.</p> <p><i>instance_hostname_SG.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

Table 4. Trace log files for troubleshooting agents (continued)

System where log is located	File name and path	Description
Definitions of variables:		
<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:SG is ON-LINE.
.
.
(42C3079B.0000-6A4:kpxreqhb.cpp,644,"HeartbeatInserter") Remote node SERVER5B:SG is OFF-LINE.
```

See the following key points about the preceding excerpts:

- The monitoring server appends the **SG** product code to the server name to form a unique name (SERVER5B:SG) for this instance of the IBM Monitoring Agent for OpenStack. 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 Monitoring Agent for OpenStack, 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 78 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 Monitoring Agent for OpenStack uses RAS1 tracing and generates the logs described in Table 4 on page 79. 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:ksg ALL)
 - Maximum error tracing. KBB_RAS1=ERROR (UNIT:ksg 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 4 on page 79 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 Monitoring Agent for OpenStack uses RAS1 tracing and generates the logs described in Table 4 on page 79. The default RAS1 trace level is ERROR.

Procedure

1. Open the trace options file:
 - **UNIX systems:**
`install_dir /config/sg_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:ksg ALL) (UNIT:kra ALL)`
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 4 on page 79 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 `F1`.

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.

```
kbbcrc1.c, 400, May 29 2007, 12:54:43, 1.1, *
kbbcrcn1.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
kbbac1c.c, 400, May 29 2007, 12:54:27, 1.4, ACF1
```

```

kbbaccli.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
kdhsgnh.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 Monitoring Agent for OpenStack, 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.

Note: You can resolve some problems by ensuring that your system matches the system requirements listed in “Prerequisites” on page 5.

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

Table 5. 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 KSGCMA 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 5. 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#acpininstall).
<p>The Monitoring Agent for OpenStack is installed on a stand-alone server with OpenStack Python clients, and cannot connect to the remote OpenStack components through domain name.</p>	<p>The server where the agent is installed does not belong to the OpenStack network area.</p> <p>You must edit <code>/etc/hosts</code> file and map the domain names with IP addresses. Here are some examples:</p> <pre>9.112.250.244 ico-ds ico-ds.demo.com 9.112.250.245 ico-cs1 ico-cs1.demo.com 9.112.250.246 ico-cs2 ico-cs2.demo.com 9.112.250.247 ico-cs3 ico-cs3.demo.com 9.112.250.248 ico-vmware-region ico-vmware-region.demo.com 9.112.250.249 ico-vmware-neutron ico-vmware-neutron.demo.com 9.112.250.250 ico-kvm-region ico-kvm-region.demo.com 9.123.101.85 ico-kvm-compute ico-kvm-compute.demo.com 9.123.101.81 ico-kvm-neutron ico-kvm-neutron.demo.com</pre> <p>Note: If you use an HTTP proxy server on the agent server, you must make sure that the proxy server can access the remote OpenStack components through domain name.</p>
<p>The Monitoring Agent for OpenStack installation failed with the following error:</p> <pre>[root@rhel66 KSGinstaller]# ./install.sh -bash: ./install.sh: /bin/ksh: bad interpreter: No such file or directory [root@rhel66 KSGinstaller]# yum install ksh</pre>	<p>Ksh(KornShell) is not installed.</p> <p>Install ksh(KornShell) before you install the agent.</p>
<p>The agent installation failed with the following error:</p> <pre>runGSkit failure: Return error code: 99. runGSkit failure: GSKit check failure, script: /opt/IBM/ITM/ls3263/gs/bin/private_verifyinstall. runGSkit failure: ls3263 - GSK check error, verifyInstall test failed. install.sh failure: Return error code: 4.</pre>	<p>The environment does not have package <code>compat-libstdc++-33-3.2.3-69.el6.i686.rpm</code> and <code>compat-libstdc++-33-3.2.3-69.el6.x86_64.rpm</code>.</p> <p>Install the package <code>compat-libstdc++-33-3.2.3-69.el6.i686.rpm</code> and <code>compat-libstdc++-33-3.2.3-69.el6.x86_64.rpm</code>.</p>

Table 5. Problems and solutions for installation and configuration (continued)

Problem	Solution
On Redhat 7, if you installed the OS agent within IBM Tivoli Monitoring V6.3 FP2 by running <code>setarch \$(uname -m) --uname-2.6</code> , the installation of the Monitoring Agent for OpenStack to the same directory fails.	Follow the workaround to install the Monitoring Agent for OpenStack and the OS agent in the same directory: <ol style="list-style-type: none"> 1. Install the Monitoring Agent for OpenStack to a directory that is different from the OS agent installation directory. 2. Replace the <code>OSAgent_dir/bin/dynarch.shl</code> file with the <code>OpenStackAgent_dir/bin/dynarch.shl</code>. 3. Install the Monitoring Agent for OpenStack again to the same directory as the OS agent.
There is no Monitoring Agent for OpenStack application support installation image for non-Linux platforms.	If your TEMS or TEPS is running on Windows or UNIX platforms, please leverage self-describing feature to install Monitoring Agent for OpenStack application support automatically. See “Enabling the self-describing capability to automatically install application support” on page 13 for more information.

Table 6. General problems and solutions for uninstallation

Problem	Solution
The way to remove inactive managed systems (systems whose status is OFFLINE) from the Navigator tree in the portal is not obvious.	Use the following steps to remove, but not uninstall, an offline managed system from the Navigator tree: <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 6. 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:KSG</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 KSG, 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>

Remote deployment troubleshooting

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

Table 7 contains problems and solutions related to remote deployment.

Table 7. Remote deployment problems and solutions

Problem	Solution
<p>While you are using the remote deployment feature to install the IBM Monitoring Agent for OpenStack, 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 7. Remote deployment problems and solutions (continued)

Problem	Solution
<p>Remote installation or configuration of Monitoring Agent for OpenStack fails. The agent instance cannot be started or stopped from Tivoli Enterprise Portal.</p>	<p>The Monitoring Agent for OpenStack registers the following strings as the managed system name: <i>instancename:hostname:SG</i>, where <i>instancename</i> is the name that you assign to the agent; <i>hostname</i> is the host name. This problem occurs because the managed system name is longer than 32 characters. When the managed system name is too long, the <i>instancename:hostname</i> part is truncated.</p> <p>To avoid this problem, shorten the instance name or the host name to make sure that the managed system name does not exceed 32 characters.</p> <ol style="list-style-type: none"> 1. Open the configuration file for the monitoring agent, which is in the following path: <i>install_dir/config/sg_instance_name.config</i>. 2. Find the line that begins with <i>CTIRA_HOSTNAME=</i>. 3. Specify a unique and shorter host name for the host computer. The final concatenated name that includes the instance name, new host name, and SG, cannot exceed 32 characters. Note: Make sure that the resulting name is unique regarding any existing monitoring component that was previously registered with the Tivoli Enterprise Monitoring Server. 4. Save the file. 5. Restart the agent.
<p>The removal of a monitoring agent fails when you use the remote removal process in the Tivoli Enterprise Portal desktop or browser.</p>	<p>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.</p>

Agent troubleshooting

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

Table 8 contains problems and solutions that can occur with the agent after it is installed.

Table 8. Agent problems and solutions

Problem	Solution
<p>Log data accumulates too rapidly.</p>	<p>Check the RAS trace option settings, which are described in “Setting RAS trace parameters by using the GUI” on page 83. The trace option settings that you can set on the <i>KBB_RAS1=</i> and <i>KDC_DEBUG=</i> lines potentially generate large amounts of data.</p>

Table 8. Agent problems and solutions (continued)

Problem	Solution
<p>The Monitoring Agent for OpenStack instance node displays incorrectly Tivoli Enterprise Portal when the managed system name <i>instancename:hostname:SG</i> is too long.</p>	<p>When the managed system name <i>instancename:hostname:SG</i> is longer than 32 characters, the <i>instancename:hostname</i> part is truncated. After the truncation, if there are agent instances with the same node name, and these agent instances are connected to the same hub Tivoli Enterprise Monitoring Server, the agent instances do not work correctly.</p> <p>Shorten the host name portion of the name by completing the following steps:</p> <ol style="list-style-type: none"> 1. Open the configuration file for the monitoring agent, which is in the following path: <i>install_dir/config/sg_instance_name.config</i>. 2. Find the line that begins with <i>CTIRA_HOSTNAME=</i>. 3. Specify a unique and shorter host name for the host computer. The final concatenated name that includes the instance name, new host name, and SG, cannot exceed 32 characters. Note: Make sure that the resulting name is unique regarding any existing monitoring component that was previously registered with the Tivoli Enterprise Monitoring Server. 4. Save the file. 5. Restart the agent.
<p>There is a delay between OpenStack environment and the monitoring data in Tivoli Enterprise Portal.</p>	<p>The Monitoring Agent for OpenStack data collector collects monitoring data every 300 seconds, and sends the data to IBM Tivoli Monitoring every 300 seconds by default. A delay might occur between the OpenStack environment and the monitoring data in Tivoli Enterprise Portal. To reduce the delay, you can shorten the data collection and emitting intervals by changing the settings in the data collector configuration file to see the updated data sooner.</p> <p>Complete the following steps to shorten the intervals:</p> <ol style="list-style-type: none"> 1. Locate and edit the following agent instance data collector configuration file: <i>CANDLEHOME/\$ITM_BINARCH/sg/bin/ksg_dc_instance_name.cfg</i> 2. Edit the value of the parameter data_collect_interval to change the data collection interval. 3. Edit the value of the parameter data_emit_interval to change the data emitting interval. 4. Save the configuration file. 5. Stop the agent instance by running the command: <i>CANDLEHOME/bin/itmcmd agent -o instance_name stop sg</i> 6. Start the agent instance by running the command: <i>CANDLEHOME/bin/itmcmd agent -o instance_name start sg</i>

Table 8. Agent problems and solutions (continued)

Problem	Solution
<p>When you use the Tivoli Enterprise Portal browser client based on IBM 1.5 JRE that is provided before ITM 6.2.3 FP5, the strings of the agents are not displayed correctly.</p>	<p>Follow the steps to identify a higher JRE version to the Tivoli Enterprise Portal Server for the client to use:</p> <ol style="list-style-type: none"> 1. Go to the computer where the Tivoli Enterprise Portal Server is installed. 2. Find the file <code>ITM_install_dir/arch/cw/jrelevel.js</code>. 3. Edit the <code>jrelevel.js</code> file by changing the value of the variable <code>jreLevel</code> to 1.6.0 as in the following example: <pre>jreLevel = "1.6.0"</pre> <p>This variable has the following possible values:</p> <ul style="list-style-type: none"> * The browser client uses the default JER for the computer. <p>1.6.0 The browser client uses the IBM 1.6 JRE.</p> <p>6.0 The browser client uses the latest Sun 1.6.0_xx JRE installed.</p>
<p>OpenStack agent has repeated GSKit error 420 in the logs.</p>	<p>For details on the cause of this issue, see the ITM technote Off-line agent due to GSKit errors.</p> <p>Follow the solution that is provided in the technote Off-line agent due to GSKit errors to add <code>ip.ssl use:n</code> to the <code>KDC_FAMILIES</code> parameter in the agent configuration file. For example:</p> <ol style="list-style-type: none"> 1. Edit the <code>install_dir/sg.ini</code> file by changing <code>KDC_FAMILIES=\$NETWORKPROTOCOL\$</code> to <code>KDC_FAMILIES='ip.ssl use:n ip.pipe port:1918 ip use:n ip.spipe use:n sna use:n'</code>. 2. If you have existing agent instances, edit the <code>install_dir/config/sg_instancename.config</code> files by changing <code>export KDC_FAMILIES='ip.pipe port:1918 ip use:n ip.spipe use:n sna use:n'</code> to <code>export KDC_FAMILIES='ip.ssl use:n ip.pipe port:1918 ip use:n ip.spipe use:n sna use:n'</code>. 3. Restart the agent.
<p>When using the <code>itmcmd agent</code> commands to start or stop this monitoring agent, you receive the following error message:</p> <pre>MKCIIN0201E Specified product is not configured.</pre>	<p>Include the command option <code>-o</code> to specify the instance to start or stop. The instance name must match the name used for configuring the agent. For example:</p> <pre>./itmcmd agent -o Test1 start sg</pre> <p>For more information about using the <code>itmcmd</code> commands, see the <i>IBM Tivoli Monitoring Command Reference</i>.</p>

Table 8. 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 8. 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.)

Workspace troubleshooting

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

Table 9 contains problems and solutions related to workspaces.

Table 9. Workspace problems and solutions

Problem	Solution
The used vCPU, Memory, and File System Storage percent are greater than 100 percent in Hypervisors and Regions workspaces.	<p>OpenStack allows users to overcommit the resources (e.g., Hypervisor Storage, Memory & vCPU). So the resource usage percent may over 100%. For resource overcommitment, please refer to the topic Overcommitting (http://docs.openstack.org/openstack-ops/content/compute_nodes.html#overcommit).</p> <p>Change the thresholds of the situations that are related to resource usage according to the resource allocation ratios. For example, if the cpu_allocation_ratio is 16, you can define a situation to show the high vCPU usage when the used vCPU percentage is approaching 1600%.</p>

Table 9. Workspace problems and solutions (continued)

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 9. 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 “Managing historical data” 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 “Managing historical data” 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 “Managing historical data” 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 “Installing and enabling application support” in the <i>IBM Tivoli Monitoring Installation and Setup Guide</i>.</p>

Situation troubleshooting

Problems can occur with situations and situation configuration.

Table 10 contains problems and solutions for situations.

Table 10. 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 83. 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 66 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 10. 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 “Situations are not firing” 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 10. 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 11 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 4 on page 79.

Table 11. 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 100. 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 Monitoring Agent for OpenStack 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:

KSG General Monitoring Agent for OpenStack 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 Monitoring Agent for OpenStack.

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 Monitoring Agent for OpenStack 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 KSG_Base and is defined in the `ksg.baroc` (version 7.5) file. The KSG_Base event class can be used for generic rules processing for any event from the IBM Monitoring Agent for OpenStack.

For events that are generated by situations in the API Endpoint Details attribute group, events are sent by using the ITM_KSG_API_DETAILS event class. This event class contains the following slots:

- `api_detecting_failure_times`: INTEGER
- `api_detecting_failure_times_enum`: STRING
- `api_detecting_times`: INTEGER
- `api_detecting_times_enum`: STRING
- `api_failure_percent`: INTEGER
- `api_failure_percent_enum`: STRING
- `api_status_code`: INTEGER
- `api_status_code_enum`: STRING
- `component_name`: STRING
- `keystone_host`: STRING
- `ksg_status`: INTEGER
- `ksg_status_enum`: STRING
- `node`: STRING
- `public_url`: STRING
- `region_name`: STRING
- `timestamp`: STRING

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

- available_api_endpoints: INTEGER
- available_api_endpoints_enum: STRING
- component_name: STRING
- keystone_host: STRING
- node: STRING
- timestamp: STRING
- unavailable_api_endpoints: INTEGER
- unavailable_api_endpoints_enum: STRING

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

- active_instance_count: INTEGER
- active_instance_count_enum: STRING
- available_api_count: INTEGER
- available_api_count_enum: STRING
- file_system_storage_usage_critical_regions: INTEGER
- file_system_storage_usage_critical_regions_enum: STRING
- file_system_storage_usage_normal_regions: INTEGER
- file_system_storage_usage_normal_regions_enum: STRING
- file_system_storage_usage_warning_regions: INTEGER
- file_system_storage_usage_warning_regions_enum: STRING
- floating_ip_usage_critical_regions: INTEGER
- floating_ip_usage_critical_regions_enum: STRING
- floating_ip_usage_normal_regions: INTEGER
- floating_ip_usage_normal_regions_enum: STRING
- floating_ip_usage_warning_regions: INTEGER
- floating_ip_usage_warning_regions_enum: STRING
- inactive_instance_count: INTEGER
- inactive_instance_count_enum: STRING
- keystone_host: STRING
- memory_usage_critical_regions: INTEGER
- memory_usage_critical_regions_enum: STRING
- memory_usage_normal_regions: INTEGER
- memory_usage_normal_regions_enum: STRING
- memory_usage_warning_regions: INTEGER
- memory_usage_warning_regions_enum: STRING
- node: STRING
- running_process_count: INTEGER
- running_process_count_enum: STRING
- tenant_name: STRING
- timestamp: STRING
- unavailable_api_count: INTEGER
- unavailable_api_count_enum: STRING

- vcpu_usage_critical_regions: INTEGER
- vcpu_usage_critical_regions_enum: STRING
- vcpu_usage_normal_regions: INTEGER
- vcpu_usage_normal_regions_enum: STRING
- vcpu_usage_warning_regions: INTEGER
- vcpu_usage_warning_regions_enum: STRING

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

- ephemeral_disk: INTEGER
- ephemeral_disk_enum: STRING
- flavor_id: STRING
- flavor_name: STRING
- is_public: STRING
- keystone_host: STRING
- memory: INTEGER
- memory_enum: STRING
- node: STRING
- reference_count: INTEGER
- reference_count_enum: STRING
- region_name: STRING
- root_disk: INTEGER
- root_disk_enum: STRING
- swap_disk: INTEGER
- swap_disk_enum: STRING
- timestamp: STRING
- vcpu_count: INTEGER
- vcpu_count_enum: STRING

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

- cpu_architecture: STRING
- cpu_model: STRING
- cpu_vendor: STRING
- host_ip: STRING
- hypervisor_hostname: STRING
- hypervisor_type: STRING
- hypervisor_version: STRING
- id: STRING
- keystone_host: STRING
- local_storage_size: STRING
- node: STRING
- region_name: STRING
- running_vms: INTEGER
- running_vms_enum: STRING
- service_host: STRING

- timestamp: STRING
- total_memory_size: STRING
- total_vcpu_count: INTEGER
- total_vcpu_count_enum: STRING
- used_memory_percent: INTEGER
- used_memory_percent_enum: STRING
- used_memory_size: STRING
- used_storage_percent: INTEGER
- used_storage_percent_enum: STRING
- used_storage_size: STRING
- used_vcpu_count: INTEGER
- used_vcpu_count_enum: STRING
- used_vcpu_percent: INTEGER
- used_vcpu_percent_enum: STRING

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

- accept_ranges: STRING
- container_name: STRING
- content_type: STRING
- keystone_host: STRING
- node: STRING
- object_count: INTEGER
- object_count_enum: STRING
- tenant_name: STRING
- timestamp: STRING
- used_bytes: REAL
- used_bytes_enum: STRING

For events that are generated by situations in the Object Storage Account Statistics attribute group, events are sent by using the ITM_KSG_ACCOUNT_STATISTICS event class. This event class contains the following slots:

- accept_ranges: STRING
- container_count: INTEGER
- container_count_enum: STRING
- content_type: STRING
- keystone_host: STRING
- node: STRING
- object_count: INTEGER
- object_count_enum: STRING
- tenant_name: STRING
- timestamp: STRING
- used_bytes: REAL
- used_bytes_enum: STRING

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

- bytes: REAL
- bytes_enum: STRING
- container_name: STRING
- content_type: STRING
- keystone_host: STRING
- last_modified: STRING
- node: STRING
- object_name: STRING
- tenant_name: STRING
- timestamp: STRING

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

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

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

- component_name: STRING
- cpu: REAL
- cpu_enum: STRING
- host_ip: STRING
- keystone_host: STRING
- ksg_hostname: STRING

- ksg_status: INTEGER
- ksg_status_enum: STRING
- mem: REAL
- mem_enum: STRING
- node: STRING
- pid: INTEGER
- pid_enum: STRING
- ppid: INTEGER
- ppid_enum: STRING
- process_name: STRING
- start_time: STRING
- timestamp: STRING

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

- component_name: STRING
- component_process_count: INTEGER
- component_process_count_enum: STRING
- keystone_host: STRING
- node: STRING
- timestamp: STRING

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

- host_process_count: INTEGER
- host_process_count_enum: STRING
- keystone_host: STRING
- ksg_hostname: STRING
- node: STRING
- timestamp: STRING

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

- active_instance_count: INTEGER
- active_instance_count_enum: STRING
- allocated_floating_ip_count: INTEGER
- allocated_floating_ip_count_enum: STRING
- allocated_floating_ip_percent: INTEGER
- allocated_floating_ip_percent_enum: STRING
- floating_ip_quota: INTEGER
- floating_ip_quota_enum: STRING
- inactive_instance_count: INTEGER
- inactive_instance_count_enum: STRING
- keystone_host: STRING
- node: STRING
- region_name: STRING

- timestamp: STRING
- total_file_system_storage_size: STRING
- total_floating_ip_count: INTEGER
- total_floating_ip_count_enum: STRING
- total_instance_count: INTEGER
- total_instance_count_enum: STRING
- total_memory_size: STRING
- total_vcpu_count: INTEGER
- total_vcpu_count_enum: STRING
- used_file_system_storage_percent: INTEGER
- used_file_system_storage_percent_enum: STRING
- used_file_system_storage_size: STRING
- used_memory_percent: INTEGER
- used_memory_percent_enum: STRING
- used_memory_size: STRING
- used_vcpu_count: INTEGER
- used_vcpu_count_enum: STRING
- used_vcpu_percent: INTEGER
- used_vcpu_percent_enum: STRING

For events that are generated by situations in the Take Action Status attribute group, events are sent by using the ITM_KSG_TAKE_ACTION_STATUS event class. This event class contains the following slots:

- action_app_return_code: INTEGER
- action_command: STRING
- action_id: INTEGER
- action_instance: STRING
- action_message: STRING
- action_name: STRING
- action_node: STRING
- action_owner: STRING
- action_results: STRING
- action_status: INTEGER
- action_status_enum: STRING
- action_subnode: STRING
- action_type: INTEGER
- action_type_enum: STRING
- node: STRING
- timestamp: STRING

For events that are generated by situations in the Thread Pool Status attribute group, events are sent by using the ITM_KSG_THREAD_POOL_STATUS event class. This event class contains the following slots:

- node: STRING
- thread_pool_active_threads: INTEGER
- thread_pool_active_threads_enum: STRING
- thread_pool_avg_active_threads: REAL
- thread_pool_avg_active_threads_enum: STRING

- thread_pool_avg_job_wait: REAL
- thread_pool_avg_job_wait_enum: STRING
- thread_pool_avg_queue_length: REAL
- thread_pool_avg_queue_length_enum: STRING
- thread_pool_max_active_threads: INTEGER
- thread_pool_max_active_threads_enum: STRING
- thread_pool_max_queue_length: INTEGER
- thread_pool_max_queue_length_enum: STRING
- thread_pool_max_size: INTEGER
- thread_pool_max_size_enum: STRING
- thread_pool_min_active_threads: INTEGER
- thread_pool_min_active_threads_enum: STRING
- thread_pool_min_queue_length: INTEGER
- thread_pool_min_queue_length_enum: STRING
- thread_pool_queue_length: INTEGER
- thread_pool_queue_length_enum: STRING
- thread_pool_size: INTEGER
- thread_pool_size_enum: STRING
- thread_pool_total_jobs: INTEGER
- thread_pool_total_jobs_enum: STRING
- timestamp: STRING

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

- availability_zone: STRING
- created_time: STRING
- flavor: STRING
- host: STRING
- hypervisor_hostname: STRING
- image: STRING
- instance_id: STRING
- instance_name: STRING
- ip_addresses: STRING
- ip_key: STRING
- keystone_host: STRING
- key_name: STRING
- ksg_status: STRING
- launched_time: STRING
- name: STRING
- node: STRING
- region_name: STRING
- tenant_id: STRING
- tenant_name: STRING
- timestamp: STRING
- update_time: STRING

- user_email: STRING
- user_id: STRING
- user_name: STRING

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

- attachments: STRING
- availability_zone: STRING
- bootable: STRING
- encrypted: STRING
- host: STRING
- keystone_host: STRING
- ksg_status: STRING
- metadata: STRING
- node: STRING
- project: STRING
- project_id: STRING
- region_name: STRING
- size: INTEGER
- size_enum: STRING
- timestamp: STRING
- type: STRING
- volume_id: STRING
- volume_name: STRING

Appendix B. Documentation library

Various publications are relevant to the use of the IBM Monitoring Agent for OpenStack.

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 Monitoring Agent for OpenStack library

The documentation for this agent and other product components is in the IBM Tivoli Monitoring Knowledge Center (http://www-01.ibm.com/support/knowledgecenter/SSTFXA_6.3.0/com.ibm.itm.doc_6.3/welcome.htm) IBM Tivoli Monitoring for Virtual Environments Knowledge Center (http://www-01.ibm.com/support/knowledgecenter/SS9U76_7.2.0.3/com.ibm.tivoli.itmvs.doc/welcome_ve72fp3.htm) ITCAM for Microsoft Applications Knowledge Center (http://www-01.ibm.com/support/knowledgecenter/SS9U76_7.1.0/com.ibm.tivoli.itmvs.doc_7.1/welcome_ve71.htm) ITCAM for Applications Knowledge Center (http://www-01.ibm.com/support/knowledgecenter/SS3JRN_7.2.0/com.ibm.itcama.doc_7.2/welcome_itcamfapps72.html).

One document is specific to the Monitoring Agent for OpenStack. The IBM Monitoring Agent for OpenStack User's Guide provides agent-specific information for configuring, using, and troubleshooting the Monitoring Agent for OpenStack.

The *Offering Guide* also provides information about installing and configuring the component products in the offering.

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 OpenStack 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*
- *IBM Tivoli Monitoring Universal Agent User's Guide*
- *IBM Tivoli Universal Agent API and Command Programming Reference Guide*
- *IBM Tivoli Monitoring: i5/OS™ Agent User's 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: UNIX Logs OS Agent User's*
 - *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 US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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

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 Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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.

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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

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

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. 2015.

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

- Accept Ranges attribute 42, 43
- Action App Return Code attribute 56
- Action Command attribute 56
- Action ID attribute 56
- Action Instance attribute 57
- Action Message attribute 57
- Action Name attribute 57
- Action Node attribute 57
- Action Owner attribute 57
- Action Results attribute 57
- Action Status attribute 57
- Action Subnode attribute 58
- Action Type attribute 58
- Active Instances 34, 53
- activities 75
- additional information
 - attributes 29
 - situations 69
 - Take Action commands 73
 - Workspaces 23
- agent
 - functions 1
 - problems and workarounds 93
- agent configuration 16
- Agent Management Services 3
- Allocated Floating IP Count attribute 54
- API Endpoint Details attribute group 31
- API Endpoint Detecting Failure Percent attribute 31
- API Endpoint Detecting Failure Times attribute 31
- API Endpoint Detecting Status Code attribute 31
- API Endpoint Detecting Times attribute 31
- API Endpoints attribute group 33
- API Endpoints Details workspace 24
- API Endpoints workspace 24
- Attachments attribute 64
- attribute group 31
- attribute groups
 - API Endpoint Details 31
 - API Endpoints 33
 - Environment Summary 34
 - Flavors 37
 - Hypervisors 38
 - list of all 29
 - Object Containers 42
 - Object Storage Account Statistics 43
 - Objects 44
 - overview 29
 - Performance Object Status 46
 - Process Details 49
 - Processes 51
 - Processes by Host 52
 - Regions 53
 - Take Action Status 56
 - Thread Pool Status 58
 - VM Details 61
 - Volumes 64
- attributes 31
 - Accept Ranges 42, 43
 - Action App Return Code 56
 - Action Command 56
- attributes (*continued*)
 - Action ID 56
 - Action Instance 57
 - Action Message 57
 - Action Name 57
 - Action Node 57
 - Action Owner 57
 - Action Results 57
 - Action Status 57
 - Action Subnode 58
 - Action Type 58
 - Active Instances 34, 53
 - additional information 29
 - Allocated Floating IP Count 54
 - API Endpoint Details 31
 - API Endpoint Detecting Failure Percent 31
 - API Endpoint Detecting Failure Times 31
 - API Endpoint Detecting Status Code 31
 - API Endpoint Detecting Times 31
 - API Endpoints 33
 - Attachments 64
 - Availability Zone 61, 64
 - Available API Endpoints 33, 34
 - Average Collection Duration 46
 - Bootable 64
 - Bytes 44
 - Cache Hit Percent 46
 - Cache Hits 46
 - Cache Misses 46
 - Component Name 49, 51
 - Component Process Count 52
 - Container Count 43
 - Container Name 42, 45
 - Content Type 42, 43, 45
 - CPU Architecture 39
 - CPU Model 39
 - CPU Usage(%) 49
 - CPU Vendor 39
 - Created Time 61
 - Encrypted 64
 - Environment Summary 34
 - Ephemeral Disk (GB) 37
 - Error Code 46
 - Flavor 61
 - Flavor ID 37
 - Flavor Name 37
 - Flavors 37
 - Floating IP Quota 54
 - Host 61, 65
 - Host IP 39, 50
 - Host Name 50, 52
 - Host Process Count 53
 - Hypervisor Hostname 39, 61
 - Hypervisor ID 39
 - Hypervisor Type 39
 - Hypervisor Version 39
 - Hypervisors 38
 - Image 61
 - Inactive Instances 34, 54
 - Instance ID 62
 - Instance Name 62

attributes (continued)

- Intervals Skipped 47
- IP Addresses 62
- IP Key 62
- Is Public 37
- Key Name 62
- Keystone Host 32, 33, 34, 37, 40, 42, 43, 45, 50, 52, 53, 54, 62, 65
- Last Collection Duration 47
- Last Collection Finished 47
- Last Collection Start 48
- Last Modified 45
- Launched Time 62
- Local Storage Size (TB) 40
- Memory (MB) 37
- Memory Usage(%) 50
- Metadata 65
- Name 62
- Node 32, 33, 34, 37, 40, 42, 44, 45, 48, 50, 52, 53, 54, 58, 63, 65
- Number of Collections 48
- Object Containers 42
- Object Count 42, 44
- Object Name 45, 48
- Object Status 48
- Object Storage Account Statistics 43
- Object Type 48
- Objects 44
- overview 29
- Performance Object Status 46
- pid 50
- ppid 50
- Process Details 49
- Process Name 51
- Processes 51
- Processes by Host 52
- Project 65
- Project ID 65
- Public URL 32
- Query Name 49
- Reference Count 38
- Refresh Interval 49
- Region Count (Critical File System Storage Usage) 34
- Region Count (Critical Floating IP Quota Usage) 35
- Region Count (Critical Memory Usage) 35
- Region Count (Critical vCPU Usage) 35
- Region Count (Normal File System Storage Usage) 35
- Region Count (Normal Floating IP Quota Usage) 35
- Region Count (Normal Memory Usage) 35
- Region Count (Normal vCPU Usage) 35
- Region Count (Warning File System Storage Usage) 35
- Region Count (Warning Floating IP Quota Usage) 36
- Region Count (Warning Memory Usage) 36
- Region Count (Warning vCPU Usage) 36
- Region Name 32, 38, 40, 54, 63, 65
- Regions 53
- Root Disk (GB) 38
- Running Processes 36
- Running VMs 40
- Service Host 40
- Service Type 32, 33
- Size (GB) 65
- SSH Connection Status 53
- Start Time 51
- Status 32, 51, 63, 66
- Swap Disk (MB) 38
- Take Action Status 56

attributes (continued)

- Tenant 36, 42, 44, 45, 63
- Tenant ID 63
- Thread Pool Active Threads 58
- Thread Pool Avg Active Threads 59
- Thread Pool Avg Job Wait 59
- Thread Pool Avg Queue Length 59
- Thread Pool Max Active Threads 59
- Thread Pool Max Queue Length 59
- Thread Pool Max Size 59
- Thread Pool Min Active Threads 60
- Thread Pool Min Queue Length 60
- Thread Pool Queue Length 60
- Thread Pool Size 60
- Thread Pool Status 58
- Thread Pool Total Jobs 60
- Timestamp 32, 33, 36, 38, 40, 43, 44, 45, 49, 51, 52, 53, 54, 58, 60, 63, 66
- Total File System Storage Size (TB) 54
- Total Floating IP Count 55
- Total Instances 55
- Total Memory Size (GB) 40, 55
- Total vCPU Count 41
- Total vCPUs 55
- Type 66
- Unavailable API Endpoints 33, 36
- Update Time 63
- Used Bytes 43, 44
- Used File System Storage Percent 55
- Used File System Storage Size (TB) 55
- Used Floating IP Quota Percent 55
- Used Memory Percent 41, 55
- Used Memory Size (GB) 41, 56
- Used Storage Percent 41
- Used Storage Size (TB) 41
- Used vCPU Count 41
- Used vCPU Percent 41, 56
- Used vCPUs 56
- User 63
- User Email 64
- User ID 64
- vCPU Count 38
- VM Details 61
- Volume ID 66
- Volume Name 66
- Volumes 64

Availability Zone attribute 61, 64

Available API Endpoints attribute 33, 34

Average Collection Duration attribute 46

B

Bootable attribute 64

Bytes attribute 44

C

Cache Hit Percent attribute 46

Cache Hits attribute 46

Cache Misses attribute 46

calculate historical data disk space 66

capacity planning for historical data 66

cinfo command 11

commands

- tacmd addSystem 21

- Take Action 73

- Component Name attribute 49, 51
- Component Process Count attribute 52
- components 1
 - IBM Tivoli Monitoring 1
- configuration 9, 16, 18
 - agent 5
 - fields 20
 - problems and workarounds 89
 - remote 21
 - values 20
- configure 16, 18
- configure the agent 16
- configure the data collector 18
- configuring the agent 16
- configuring the agent data collector 18
- configuring the monitoring agent 5
- Container Count attribute 43
- Container Name attribute 42, 45
- Content Type attribute 42, 43, 45
- CPU Architecture attribute 39
- CPU Model attribute 39
- CPU Usage(%) attribute 49
- CPU Vendor attribute 39
- Created Time attribute 61

D

- data collector configuration 18
- descriptions 70
- detailed 83
- developerWorks website 116
- disk capacity planning for historical data 66
- documentation
 - See* publications

E

- Encrypted attribute 64
- Environment Summary attribute group 34
- Environment Summary workspace 25
- Ephemeral Disk (GB) attribute 37
- Error Code attribute 46
- event
 - mapping 105

F

- Flavor attribute 61
- Flavor ID attribute 37
- Flavor Name attribute 37
- Flavors attribute group 37
- Flavors workspace 25
- Floating IP Quota attribute 54

H

- historical data
 - calculate disk space 66
 - disk capacity planning 66
- Host attribute 61, 65
- Host IP attribute 39, 50
- Host Name attribute 50, 52
- Host Process Count attribute 53
- Hypervisor Hostname attribute 39, 61
- Hypervisor ID attribute 39

- Hypervisor Type attribute 39
- Hypervisor Version attribute 39
- Hypervisors attribute group 38
- Hypervisors workspace 25

I

- IBM Tivoli Monitoring 1
 - overview 1
- Image attribute 61
- Inactive Instances attribute 34, 54
- install 9
- install self-describing agent 13
- installation 9
 - agent 5
 - problems and workarounds 89
 - remote 21
- installing language packs 6
- installing the monitoring agent 5
- Instance ID attribute 62
- Instance Name attribute 62
- Integrated Service Management Library documentation 116
- interface
 - user 3
- Intervals Skipped attribute 47
- IP Addresses attribute 62
- IP Key attribute 62
- Is Public attribute 37

K

- Key Name attribute 62
- Keystone Host attribute 32, 33, 34, 37, 40, 42, 43, 45, 50, 52, 53, 54, 62, 65

L

- language packs 6
 - installing 6
 - silent installation 6
- Last Collection Duration attribute 47
- Last Collection Finished attribute 47
- Last Collection Start attribute 48
- Last Modified attribute 45
- Launched Time attribute 62
- list of messages 104
- Local Storage Size (TB) attribute 40

M

- Memory (MB) attribute 37
- Memory Usage(%) attribute 50
- messages
 - contents 103
 - for IBM Monitoring Agent for OpenStack 104
 - format 103
- Metadata attribute 65
- Monitoring Agent for OpenStack
 - performance considerations 100

N

- Name attribute 62
- Node attribute 32, 33, 34, 37, 40, 42, 44, 45, 48, 50, 52, 53, 54, 58, 63, 65

Number of Collections attribute 48

O

Object Containers attribute group 42
Object Containers workspace 26
Object Count attribute 42, 44
Object Name attribute 45, 48
Object Status attribute 48
Object Storage Account Statistics attribute group 43
Object Storage Account Statistics workspace 26
Object Type attribute 48
Objects attribute group 44
Objects workspace 26
OpenStack
 situations 73
 workspaces
 descriptions 24
OpenStack workspace 24
overview
 IBM Tivoli Monitoring 1

P

performance considerations 100
Performance Object Status attribute group 46
pid attribute 50
policies 75
ppid attribute 50
prerequisite publications 115
prerequisite, software 5
problems and workarounds 88
 agent-specific 93
 agent-specific workspaces 97
 configuration 89
 install 89
 remote deployment 92
 situations 100
 Take Action commands 102
 workspaces 97
Process Details attribute group 49
Process Details workspace 26
Process Name attribute 51
Processes attribute group 51
Processes by Host attribute group 52
Processes workspace 26
proddsc.tbl 11
Project attribute 65
Project ID attribute 65
Public URL attribute 32
publications 115, 116
 developerWorks website 116
 IBM Tivoli Monitoring 115
 Integrated Service Management Library 116
 prerequisite 115
 Redbooks 116
 related 116
 Technotes 116
 wikis 116

Q

queries, using attributes 29
Query Name attribute 49

R

ras1 85
Redbooks 116
Reference Count attribute 38
Refresh Interval attribute 49
Region Count (Critical File System Storage Usage) attribute 34
Region Count (Critical Floating IP Quota Usage) attribute 35
Region Count (Critical Memory Usage) attribute 35
Region Count (Critical vCPU Usage) attribute 35
Region Count (Normal File System Storage Usage) attribute 35
Region Count (Normal Floating IP Quota Usage) attribute 35
Region Count (Normal Memory Usage) attribute 35
Region Count (Normal vCPU Usage) attribute 35
Region Count (Warning File System Storage Usage) attribute 35
Region Count (Warning Floating IP Quota Usage) attribute 36
Region Count (Warning Memory Usage) attribute 36
Region Count (Warning vCPU Usage) attribute 36
Region Name attribute 32, 38, 40, 54, 63, 65
Regions attribute group 53
Regions workspace 27
remote
 installation and configuration 21
remote deployment
 problems and workarounds 92
requirements 5
response file template 6
Root Disk (GB) attribute 38
Running Processes attribute 36
Running VMs attribute 40

S

self-describing agent installation 13
Service Host attribute 40
Service Type attribute 32, 33
silent installation 6
silent installation of language packs 6
situations 70
 additional information
 predefined, defined 69
 overview 69
 predefined 69
 problems and workarounds 100
 Situation Editor 69
situations, using attributes 29
Size (GB) attribute 65
Software prerequisite 5
SSH connection 18
SSH Connection Status attribute 53
Start Time attribute 51
Status attribute 32, 51, 63, 66
support 14
 list of messages 104
support files 14
Swap Disk (MB) attribute 38

T

tacmd addSystem command 21
Take Action commands
 additional information 73
 overview 73

- Take Action commands (*continued*)
 - predefined 73, 75
 - problems and workarounds 102
- Take Action Status attribute group 56
- Technotes 116
- Tenant attribute 36, 42, 44, 45, 63
- Tenant ID attribute 63
- Thread Pool Active Threads attribute 58
- Thread Pool Avg Active Threads attribute 59
- Thread Pool Avg Job Wait attribute 59
- Thread Pool Avg Queue Length attribute 59
- Thread Pool Max Active Threads attribute 59
- Thread Pool Max Queue Length attribute 59
- Thread Pool Max Size attribute 59
- Thread Pool Min Active Threads attribute 60
- Thread Pool Min Queue Length attribute 60
- Thread Pool Queue Length attribute 60
- Thread Pool Size attribute 60
- Thread Pool Status attribute group 58
- Thread Pool Total Jobs attribute 60
- Timestamp attribute 32, 33, 36, 38, 40, 43, 44, 45, 49, 51, 52, 53, 54, 58, 60, 63, 66
- Tivoli Enterprise Console
 - event mapping 105
- Tivoli Enterprise Monitoring Server support 14
- Tivoli Enterprise Portal desktop client support 16
- Tivoli Enterprise Portal Server support 15
- Total File System Storage Size (TB) attribute 54
- Total Floating IP Count attribute 55
- Total Instances attribute 55
- Total Memory Size (GB) attribute 40, 55
- Total vCPU Count attribute 41
- Total vCPUs attribute 55
- trace
 - turn off 87
 - turn on 87
- trace settings 85
- tracing 83
- troubleshooting 77
 - agent-specific 93
 - agent-specific workspaces 97
 - installation 89
 - problems and workarounds 88
 - remote deployment 92
 - situations 100
 - Take Action commands 102
 - turn off trace 87
 - turn on trace 87
 - uninstallation 89
 - workspaces 97
- Type attribute 66

U

- Unavailable API Endpoints attribute 33, 36
- Update Time attribute 63
- Used Bytes attribute 43, 44
- Used File System Storage Percent attribute 55
- Used File System Storage Size (TB) attribute 55
- Used Floating IP Quota Percent attribute 55
- Used Memory Percent attribute 41, 55
- Used Memory Size (GB) attribute 41, 56
- Used Storage Percent attribute 41
- Used Storage Size (TB) attribute 41
- Used vCPU Count attribute 41
- Used vCPU Percent attribute 41, 56
- Used vCPUs attribute 56

- User attribute 63
- User Email attribute 64
- User ID attribute 64
- user interface options 3

V

- vCPU Count attribute 38
- views
 - API Endpoints Details workspace 24
 - API Endpoints workspace 24
 - Environment Summary workspace 25
 - Flavors workspace 25
 - Hypervisors workspace 25
 - Object Containers workspace 26
 - Object Storage Account Statistics workspace 26
 - Objects workspace 26
 - OpenStack workspace 24
 - Process Details workspace 26
 - Processes workspace 26
 - Regions workspace 27
 - Virtual Machines workspace 28
 - Volumes workspace 28
- Virtual Machines workspace 28
- VM Details attribute group 61
- Volume ID attribute 66
- Volume Name attribute 66
- Volumes attribute group 64
- Volumes workspace 28

W

- wikis 116
- workarounds 88
- Workflow Editor 75
- workspaces
 - API Endpoints 24
 - API Endpoints Details 24
 - descriptions 24
 - Environment Summary 25
 - Flavors 25
 - Hypervisors 25
 - Object Containers 26
 - Object Storage Account Statistics 26
 - Objects 26
 - OpenStack 24
 - predefined 23
 - problems and workarounds 97
 - Process Details 26
 - Processes 26
 - Regions 27
 - Virtual Machines 28
 - Volumes 28
- Workspaces
 - additional information 23
 - overview 23



Printed in USA

SC27-6586-01

