IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage

Version 7.3.0

User's Guide

SC14-7489-02

IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage

Version 7.3.0

User's Guide

SC14-7489-02

Note
Before using this information and the product it supports, read the information in "Notices" on page 105.

This edition applies to version 7.3.0 of IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage (product number 5724-L92) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Tables

Chapter 1. Overview of the agent

The IBM® Tivoli® Monitoring for Virtual Environments Agent for NetApp Storage provides you with the capability to monitor NetApp and IBM N Series storage systems through OnCommand API Services. IBM Tivoli Monitoring is the base software for the NetApp Storage agent.

IBM Tivoli Monitoring overview

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

Functions of the monitoring agent

Displays health and performance of NetApp and IBM N series storage systems
: You can use the NetApp Storage agent to visualize capacity, latency, and throughput performance metrics of NetApp and IBM N series storage systems.

Helps identify poorly performing storage system objects
: The NetApp Storage agent collects several important performance metrics for four important NetApp storage objects: aggregates, volumes, and logical unit numbers (LUNs); sorts object lists by performance metrics to show the least healthy or busiest objects at the top; and displays trend graphs of the same performance metrics when showing an object in detail.

Components of the IBM Tivoli Monitoring environment

After you install and set up the NetApp Storage agent, you have an environment that contains the client, server, and monitoring agent implementation for IBM Tivoli Monitoring.

This IBM 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, NetApp Storage agent (one or more instances of the monitoring agent) The instances communicate with the systems or subsystems that you want to monitor. This monitoring agent collects and distributes data to a Tivoli Enterprise Portal Server.

IBM Tivoli Enterprise Console®
The Tivoli Enterprise Console is an optional component that acts as a central collection point for events from a variety of 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 IBM Tivoli Monitoring situations to the Tivoli Enterprise Console component.

IBM Tivoli Netcool/OMNIbus
Tivoli Netcool/OMNIbus is an optional component and an alternative to the Tivoli Enterprise Console. The Netcool/OMNIbus software is a service level management (SLM) system that delivers real-time, centralized monitoring of complex networks and IT domains. The Tivoli Netcool/OMNIbus components work together to collect and manage network event information.

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. For IBM Tivoli Monitoring for Virtual Environments, you can use Tivoli Common Reporting as a separate installation or as part of the IBM Tivoli Monitoring for Virtual Environments Performance and Capacity Management Reports capability.

IBM Tivoli Monitoring for Virtual Environments Dashboard, reporting, and Capacity Planner capabilities
The dashboard capability provides a summary view of the health of the entire environment so you can quickly assess if a problem exists and take action to address the problem. Predefined performance and capacity management reports provide a complete assessment of the capacity (including forecast) of the virtual environment based on actual historical usage. Capacity planner analytics and reports provide the ability to create what-if planning scenarios that can be used to optimize and consolidate the virtual environment.

Agent Management Services

You can use IBM Tivoli Monitoring Agent Management Services to manage the NetApp Storage agent.

Agent Management Services is available for the following IBM Tivoli Monitoring OS agents: Windows, Linux, and UNIX. The services are designed to keep the NetApp Storage agent available, and to provide information about the status of the product to the Tivoli Enterprise Portal. IBM Tivoli Monitoring V6.3.0, Fix Pack 2 or later provides support for Agent Management Services. For more information about Agent Management Services, see the *IBM Tivoli Monitoring Administrator's Guide*, "Agent Management Services" chapter.

User interface options

Installation of the base Tivoli monitoring software and other integrated applications provides a variety of 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. You can use Internet Explorer to access the browser client.

Command-line interface

You can use IBM 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 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.

IBM Tivoli Netcool/OMNIbus event list

You can use the event list to monitor and manage alerts. An alert is created when the ObjectServer receives an event, alarm, message, or data item. Each alert is made up of columns (or fields) of information that are held 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.

Tivoli Common Reporting

Use the Tivoli Common Reporting web user interface when you have installed Tivoli Common Reporting separately. In this interface, you specify report parameters and other report properties, generate formatted reports, schedule reports, and view reports. This user interface is based on the Tivoli Integrated Portal.

IBM Tivoli Monitoring for Virtual Environments Dashboard, reporting, and Capacity Planner capabilities

This user interface is based on the Tivoli Integrated Portal. The Dashboard provides predefined contextual summary views of the health (availability, performance, and capacity) of the complete virtual environment. Performance and Capacity Management Reports provides predefined Cognos-based reports that contain historical data, and a data model with tools for creating ad hoc reports. Capacity Planner provides you with a tool to import data for analysis and observe trends and patterns that you use to generate recommendations and create reports in the dashboard.

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.

Before installing and configuring the agent, make sure your environment meets the requirements for the IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage. See the Prerequisites topic for the agent in the the IBM Tivoli Monitoring for Virtual Environments knowledge center.

To install and configure the NetApp Storage agent, use the procedures for installing monitoring agents in the *IBM Tivoli Monitoring Installation and Setup Guide* along with the agent-specific installation and configuration information.

If you are doing a silent installation by using a response file, see that information in the *IBM Tivoli Monitoring Installation and Setup Guide*.

Prerequisites

- To monitor the performance counters, you must configure the OnCommand API Services with the OnCommand Unified Manager (OCUM) and OnCommand Performance Manager.
- To obtain performance data from the OnCommand API Services, the configured user must be of the type "administrator" or "monitor". These user types have default permissions to execute the REST API.
- Versions of the OnCommand API Services, OCUM, and the Performance Manager must be compatible. For example, to configure API Service V 1.0, pair OCUM V 6.2, 6.1, or 6.0 with Performance Manager V 1.1.

  For compatible product versions, see the interoperability matrix tool: http://mysupport.netapp.com/matrix/#search

Note: For more details about installing and configuring the OCUM, OnCommand Performance Manager, and the OnCommand API Services, see the *Installation and Setup Guide* of the respective component.

Language pack installation

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

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

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

Installing language packs on Windows systems

You can install the language packs on a Windows system.

Before you begin

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

Procedure

1. On the language pack CD, double-click the `lpinstaller.bat` file to start the installation program.
2. Select the language of the installer and click **OK**.
3. In the Introduction panel, click **Next**
4. Click **Add/Update** and click **Next**.
5. Select the folder where the National Language Support package (`NLSPackage`) files are located. Typically, the NLSPackage files are located in the `nlspackage` folder where the installer executable file is located.
6. Select the language support for the agent of your choice and click **Next**. To make multiple selections, press Ctrl and select the language that you want.
7. Select the languages that you want to install and click **Next**.
8. Examine the installation summary page and click **Next** to begin installation.
9. After installation completes, click **Finish** to exit the installer.
10. Restart the Tivoli Enterprise Portal, Tivoli Enterprise Portal Server, and Eclipse Help Server if any of these components are installed.

Installing language packs on UNIX or Linux systems

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

Before you begin

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

Procedure

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

Silent installation of language packs for agents

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."
2. Change the following parameter settings:

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

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

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

   LANG_SELECTION_LIST
   Language you want to install.
3. Enter the command to install the language pack with a response file (silent installation):
   - For Windows systems:

     `lpinstaller.bat -f path_to_response_file`
   - For UNIX or Linux systems:

     `lpinstaller.sh -c candle_home -f path_to_response_file`

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

Response file example

```
#     IBM Tivoli Monitoring Agent Language Pack Silent Installation Operation
#
#This is a sample response file for silent installation mode for the IBM Tivoli
#Monitoring Common Language Pack Installer.
#.
#This file uses the IBM Tivoli Monitoring Common Agent Language Pack with the
#install package as an example.
#Note:
#This response file is for the INSTALLATION of language packs only.
#This file does not support UNINSTALLATION of language packs in silent mode.
#-----------------------------------------------------------------------------
#-----------------------------------------------------------------------------
#To successfully complete a silent installation of the the example of Common Agent
#localization pack, complete the following steps:
#
#1.Copy ITM_Agent_LP_silent.rsp to the directory where lpinstaller.bat or
#lpinstaller.sh is located (IBM Tivoli Monitoring Agent Language Pack build
#location).
#
#2.Modify the response file so that it is customized correctly and completely for
#your site.
#  Complete all steps listed as follows 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,
#//installtivoli//lpsilenttest//nlspackage).
#-------------------------------------------------------------------------------
#NLS_PACKAGE_FOLDER=C:\\zosgmv\\LCD7-3583-01\\nlspackage
NLS_PACKAGE_FOLDER=//tmp//LP//nlspackage
#-------------------------------------------------------------------------------
#List the packages to process; both variables are required.
#Each variable requires that full paths are specified.
#Separate multiple entries with a semicolon (;).
#For Windows:
#           Use the backslash-backslash(\\) as a file separator.
#For Unix and Linux:
#           Use the slash-slash (//) as a file separator.
#-------------------------------------------------------------------------------
#PROD_SELECTION_PKG=C:\\zosgmv\\LCD7-3583-01\\nlspackage\\KIP_NLS.nlspkg
#BASE_AGENT_FOUND_PKG_LIST=C:\\zosgmv\\LCD7-3583-01\\nlspackage\\KIP_NLS.nlspkg
PROD_SELECTION_PKG=//tmp//LP//nlspackage//kex_nls.nlspkg;//tmp//LP//nlspackage//
koq_nls.nlspkg
BASE_AGENT_FOUND_PKG_LIST=//tmp//LP//nlspackage//kex_nls.nlspkg;//
tmp//LP//nlspackage//koq_nls.nlspkg
#-------------------------------------------------------------------------------
#List the languages to process.
#Separate multiple entries with semicolons.
#-------------------------------------------------------------------------------
LANG_SELECTION_LIST=pt_BR;fr;de;it;ja;ko;zh_CN;es;zh_TW
```

Agent-specific installation and configuration

In addition to the installation and configuration information in the *IBM Tivoli Monitoring Installation and Setup Guide*, use the agent-specific installation and configuration information to install NetApp Storage agent. This NetApp installation is new and you cannot upgrade it from the previous release.

User ID

The monitoring agent requires a user ID with a minimum privilege of GlobalRead on the NetApp data source that is being monitored. The user ID is used by the NetApp Storage agent to communicate with the NetApp OCUM Server.

Create a user ID in your NetApp storage system. For more information on how to create the user ID in your NetApp storage system, see the NetApp documentation.

Number of instances

Additional instances of this monitoring agent increase the load on the OCUM server and provide redundant information.

Keep the number of instances of the NetApp Storage agent that you configure for the same data source to a minimum, preferably one.

NetApp Storage agent application support

All agents require that you install application support files containing agent-specific information on the monitoring server, portal server, and portal desktop client.

See the *IBM Tivoli Monitoring Installation and Setup Guide* for information about installing application support.

Configuration values

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

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

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

Data Provider (DATA_PROVIDER)

Maximum Number Of Data Provider Log Files (KNU_LOG_FILE_MAX_COUNT)
The number of log files the data provider produces before overwriting previous log files. The default value is 10.

Maximum Size in KB of Each Data Provider Log (KNU_LOG_FILE_MAX_SIZE)
The maximum amount of data (in kilobytes) that the data provider writes to a single log file before creating a new log file. The default value is 5190.

Level of Detail in Data Provider Log. (KNU_LOG_LEVEL)
The amount of detail that the data provider includes in its log files. The default value is Info. Log levels include the following log messages:

- **Off**: No messages are logged.
- **Severe**: Only errors are logged.
- **Warning**: Everything that is logged at the Severe level and potential errors that might result in undesirable behavior.
- **Info**: Everything that is logged at the Warning level and high-level informational messages that describe the state of the data provider as it executes.
- **Fine**: Everything that is logged at the Info level and low-level informational messages that describe the state of the data provider as it executes.
- **Finer**: Everything that is logged at the Fine level plus highly-detailed informational messages, such as performance profiling information and debug data. Choosing this option can adversely affect the performance of the monitoring agent. This setting is intended only as a tool for problem determination in conjunction with IBM support staff.
- **Finest**: Everything that is logged at the Fine level and the most detailed informational messages, including low-level programming messages and data. Choosing this option might adversely affect the performance of the monitoring agent. This setting is intended only as a tool for problem determination in conjunction with IBM support staff.
- **All**: All messages are logged.

OnCommand API Service (Data Source)

Server (KNU_API_SERVICES_HOST_ADDRESS)
The host name or IP address of the OnCommand API Services that will be monitored.

User (KNU_API_SERVICES_USERNAME)

A user name to connect to OnCommand API Services. To obtain data from the OnCommand API Services, the configured user must be of the type "administrator" or "monitor". These user types have default permissions to execute the rest API.

Password (KNU_API_SERVICES_PASSWORD)
The password for the user name in this section.

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 9.

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.

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

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 NU -n OSAgentManagedSystemName -p INSTANCE=InstanceName \
    KNU_LOG_FILE_MAX_COUNT=LogFileMaxCount \
    KNU_LOG_FILE_MAX_SIZE=LogFileMaxSize \
    KNU_LOG_LEVEL=LogLevel \
    KNU_DATASOURCE_HOST_ADDRESS=HostAddress \
    KNU_DATASOURCE_PROTOCOL=Protocol \
    KNU_DATASOURCE_USERNAME=Username \
    KNU_DATASOURCE_PASSWORD=Password
```

In this command, the fields are defined as follows. For more information, see "Configuration values" on page 9.

OSAgentManagedSystemName
The managed system name of the OS agent that is running on the system to which the NetApp Storage agent is to be remotely deployed.

InstanceName
The name of the instance being deployed.

LogFileMaxCount
The maximum number of data provider log files. Valid values are positive integers.

LogFileMaxSize
The maximum size (in kilobytes) of each data provider log. Valid values are positive integers.

LogLevel
The level of detail in data provider logs. Valid values are `OFF`, `SEVERE`, `WARNING`, `INFO`, `FINE`, `FINER`, `FINEST`, and `ALL`.

HostAddress
The host name or IP address of the OnCommand Service API that is to be monitored.

Username
A user name on the OnCommand Service API.

Password
The password for the user name in *Username*.

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 19.

Predefined workspaces

The NetApp Storage agent provides predefined workspaces, which are organized by Navigator item.

- Aggregates Navigator item
  - Aggregates
  - Aggregate Details
  - Aggregate Details By Event
- Disks Navigator item
  - Disks
  - Disks Details
  - Disks Details By Event

- LUNs Navigator item
  - LUNs
- Qtrees Navigator item
  - Qtrees
- Volumes Navigator item
  - Volumes
  - Volume Details
  - Volume Details By Event
- Monitored DataSource Navigator item
  - Monitored DataSource
  - Cluster Details
- Events Navigator item
  - Events

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.

Aggregates Navigator item

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

Aggregate Details workspace

This workspace shows the details of specific aggregates with performance counters, occurred events with graphical real-time charts.

This workspace contains the following views:

Aggregate Details

This view shows all the key attribute along with the performance metrics of aggregate in tabular format.

Events

This view shows all the events occurred on selected aggregate.

I/O operations (per second)

This view shows a bar chart of all the operation related to I/O in the environment.

Navigator

This view helps user to get the details of each of the aggregate from the same page without navigating back to the parent workspace.

Latency (micro seconds)

This view shows latency related data in the form of bar chart .

Storage (GB)

This view depicts storage view of specific aggregate in the form of pie chart.

Aggregate Details By Event workspace

This workspace shows the details of specific aggregates based on the event from which user selected to navigate to it's details, it's design is same as that of aggregate details.

This workspace contains the following views:

Aggregate Details

This view shows all the key attribute along with the performance metrics of aggregate in tabular format.

Events

This view shows all the events occurred on the selected aggregate.

I/O operations (per second)
This view shows a bar chart of all the operation related to I/O in the environment.

Navigator
This view helps user to get the details of each of the aggregate based on event from the same page without navigating back to the parent workspace.

Latency (micro seconds)
This view shows latency related data in the form of bar chart .

Storage (GB)
This view depicts storage view of specific aggregate in the form of pie chart.

Aggregates workspace
This workspace shows the summary of aggregates along with list of events occurred on all the aggregates.

This workspace contains the following views:

Top 5 by Committed Space (GB)
This view shows a bar chart of the Committed Space for the aggregate. Top 5 aggregates that committed maximum space across the available aggregates gets highlighted here.

Bottom 5 by Throughput (bytes per second)
This view shows a bar chart of Throughput for the aggregate. 5 Aggregates that have Least Throughput will gets highlighted in this view.

Aggregate's Summary
This view shows summarized static data of all the aggregates.Each row of aggregate has a link to know the aggregate details.

Events
This view shows the all the events that occurred on the available aggregates across the environment.Each row of the event has a link to get details of the aggregate based on the selected event.

Top 5 By Utilization (percentage)
This view shows the bar chart of top 5 aggregates that have maximum utilization.

Disks Navigator item

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

Disk Details workspace
This workspace shows the details of specific disk with performance counters, occurred events with graphical real-time charts.

This workspace contains the following views:

Disk Details
This view shows all the configuration and performance data for selected disk in tabular format.

Latency (microsecond)
This view shows latency related data in the form of bar chart.

Events
This view shows all the events that occurred on the selected volume.

Navigator
This view helps user to get the details of each disk from the same page without navigating back to the parent workspace.Each row of navigator has a link to get the details based on selected Disk.

I/O Rate
This view shows I/O operations ratio on the disk in the form of bar chart.

Disk Busy (in percentage)
This view shows a disk busy percentage of disk on circular gauge.

Disk Details By Event workspace
This workspace shows the details of specific aggregates based on the event from which user selected to navigate to it's details, it's design is same as that of aggregate details.

This workspace contains the following views:

Disk Details
This view shows all the configuration and performance data for selected disk in tabular format.

Latency (microsecond)
This view shows latency related data in the form of bar chart.

Events
This view shows all the events that occurred on the selected volume.

Navigator
This view helps user to get the details of each disk from the same page without navigating back to the parent workspace.Each row of navigator has a link to get the details of disk based on selected event.

I/O Rate
This view shows I/O operations ratio on the disk in the form of bar chart.

Disk Busy (in percentage)
This view shows a disk busy percentage of disk on circular gauge.

Disks workspace
This workspace shows the summary of disks along with list of events occurred on all the available disks.

This workspace contains the following views:

Disk Summary
This view shows summary of all the available disks in tabular format. Each row of table has a link to get the details of selected Disk.

Events
This view shows all the events that occurred on all the disks.Each row of the table has a link to get details of the disk based on the selected event.

Top 5 Busy Disks (in percentage)
This view shows bar chart that depicts top 5 disks serving highest I/O operations.

Events Navigator item

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

Events workspace
This workspace shows all the key parameters of the event that occurred in the environment.

This workspace contains the following view:

Events
This view shows all the events occurred on all the entities, Each entity has a link to the specific object it relates, to know the details of individual.

LUNs Navigator item

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

LUNs workspace
This workspace shows all the key attributes of the LUNs along with all the events occurred on it.

This workspace contains the following views:

Lun Summary
This view shows performance and configuration related attributes of the LUNs.

Events
This view shows all the events that occurred on all the LUNs.

Monitored DataSource Navigator item

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

Cluster Details workspace
This workspace shows the details of specific aggregates with performance counters, occurred events with graphical real-time charts.

This workspace contains the following views:

Cluster Node Summary
This view shows all the key static/configuration attributes of all the cluster nodes in tabular format.

Performance Details
This view shows all the performance metrics of available cluster nodes in tabular format.

Throughput (kb per second)
This view shows a bar chart representations for throughput of cluster node.

Load (percentage)
This view helps user to get the details of load on each node in the environment have, this is represented with the help of CPU Busy and Average Processor Busy attributes.

Events
This view shows all the events occurred on the cluster nodes.

Monitored DataSource workspace
This workspace shows the summary of entire NetApp environment to know the user precisely to which entity/object he should be more concerned about.

This workspace contains the following views:

Cluster Info
This view shows information of monitored cluster in tabular format.

Datasource Details
This view gives a brief version information about the monitored environment along with the connection status of an agent.

Node Utilization(percentage)
This view represents Node utilization in percent on circular gauge.

Events Summary
This view shows count of events occurred across the environment per entity based on severity of an event, it also provides navigation link if user interested to see what is happening on the selected entity.

Aggregate Utilization(percentage)
This view represents Aggregate Node utilization in percent on circular gauge.

Events
This view shows all the events occurred on the cluster.

Qtrees Navigator item

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

Qtrees workspace
This workspace shows all the key attributes of the Qtrees along with all the events occurred on it.

This workspace contains the following views:

Qtree Summary
This view shows performance as well as configuration related attributes of Qtrees. Link provided on each row of qtree will navigate user to the associated volume for selected qtree.

Events
This view shows all the events that occurred on all the Qtrees.

Volumes Navigator item

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

Volume Details workspace
This workspace shows the details of specific volume with performance counters, occurred events with graphical real-time charts.

This workspace contains the following views:

Volume Details
This view shows all the configuration related data for selected volume in tabular format.

Performance Details
This view shows all the key performance attribute for the selected volume in tabular format.

Events
This view shows all the events that occurred on the selected volume.

Navigator
This view helps user to get the details of each volume from the same page without navigating back to the parent workspace.Each row of Navigator has a link to get the details of volume based on selected volume.

Latency (in milliseconds)
This view shows latency related data in the form of bar chart.

Size Used (in percentage)
This view shows a usage of volume in percentage on circular gauge.

Volume Details By Event workspace
This workspace shows the details of specific aggregates based on the event from which user selected to navigate to it's details, it's design is same as that of aggregate details.

This workspace contains the following views:

Volume Details
This view shows all the configuration related data for selected volume in tabular format.

Performance Details
This view shows all the key performance attribute for the selected volume in tabular format.

Events
This view shows all the events that occurred on the selected volume.

Navigator
This view helps user to get the details of each volume from the same page without navigating back to the parent workspace.Each row of navigator has a link to get the details of volume based on selected event.

Latency (in milliseconds)
This view shows latency related data in the form of bar chart .

Size Used (in percentage)
This view shows a circula gauge for the usage of volume in percentage.

Volumes workspace
This workspace shows the summary of aggregates along with list of events occurred on all the aggregates.

This workspace contains the following views:

Volume Summary
This view shows summary of all the available volumes in tabular format. Each row in a table has a link to get the details of each volume.

Events
This view shows all the events occurred on all the volumes.Each row of table has a link to get the details of volume based on selected event.

Top 5 Volumes By Latency (in milliseconds)
This view shows bar chart that depicts top 5 volumes with highest latency.

Chapter 4. Attributes reference

Attributes are the application properties that are being measured and reported by the IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage.

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 NetApp Storage agent and registers an *event* if the condition is met. You are alerted to events by indicator icons that are displayed in the Navigator.

Additional information about attributes

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

For a list of the attribute groups, a list of the attributes in each attribute group, and descriptions of the attributes for this monitoring agent, see "Attribute groups for the monitoring agent" and "Attributes in each attribute group" on page 20.

Attribute groups for the monitoring agent

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

- Attribute group name: Aggregates
  - Table name: KNU02AGREG
  - Warehouse table name: KNU_AGGREGATES or KNU02AGREG
- Attribute group name: Cluster Node
  - Table name: KNU09CLSND
  - Warehouse table name: KNU_CLUSTER_NODE or KNU09CLSND
- Attribute group name: Clusters
  - Table name: KNU08CLST
  - Warehouse table name: KNU_CLUSTERS or KNU08CLST
- Attribute group name: DataSource

  - Table name: KNU10DTSR
  - Warehouse table name: KNU_DATASOURCE or KNU10DTSR
- Attribute group name: Disks
  - Table name: KNU04DISK
  - Warehouse table name: KNU_DISKS or KNU04DISK
- Attribute group name: Events
  - Table name: KNU07EVENT
  - Warehouse table name: KNU_EVENTS or KNU07EVENT
- Attribute group name: EventsOverview
  - Table name: KNU11EVOW
  - Warehouse table name: KNU_EVENTSOVERVIEW or KNU11EVOW
- Attribute group name: LUNs
  - Table name: KNU05LUN
  - Warehouse table name: KNU_LUNS or KNU05LUN
- Attribute group name: Performance Object Status
  - Table name: KNUPOBJST
  - Warehouse table name: KNU_PERFORMANCE_OBJECT_STATUS or KNUPOBJST
- Attribute group name: Qtrees
  - Table name: KNU06QTREE
  - Warehouse table name: KNU_QTREES or KNU06QTREE
- Attribute group name: Volumes
  - Table name: KNU03VOL
  - Warehouse table name: KNU_VOLUMES or KNU03VOL

Attributes in each attribute group

Attributes in each NetApp Storage agent attribute group collect data that the agent uses for monitoring.

The descriptions of the attribute groups contain the following information:

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

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

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

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

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

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

Aggregates attribute group

The Aggregates attribute group contains information about the aggregates.

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 Aggregates attribute group:

Aggregate Key attribute **This attribute is a key attribute.**

Description

The resource key for this aggregate.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

AGGREGATE_KEY or AGGR_KEY

Aggregate Name attribute **This attribute is a key attribute.**

Description

Name of the aggregate

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

AGGREGATE_NAME or AGGR_NAME

Aggregate State attribute **This attribute is a key attribute.**

Description

State of the aggregate.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

AGGREGATE_STATE or AGGR_STATE

Aggregate Type attribute

Description

Type of the aggregate.

Type

string with enumerated values. The following values are defined: unavailable (unavailable), traditional (traditional), aggregate (aggregate). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

AGGREGATE_TYPE or AGGR_TYPE

Avail Size attribute

Description

Available size (in GB) for the aggregate.

Type

real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

AVAIL_SIZE

Block Rate attribute

Description

The number of block-based read and write operations that are performed per second for this aggregate.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
BLOCK_RATE

Block Type attribute

Description
The indirect block format of the aggregate. The format can be 32_bit or 64_bit. The 64_bit value indicates that the associated aggregates can be larger than 16 TB.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
BLOCK_TYPE

Cluster Key attribute **This attribute is a key attribute.**

Description
The resource key for the cluster.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
CLUSTER_KEY or CLST_KEY

CP Read Blocks attribute

Description
The number of blocks that are read per second on this aggregate during the consistency point.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
CP_READ_BLOCKS or CP_RD_BLK

CP Read Latency attribute

Description
The average latency (in microseconds) per block for the read operations during a consistency point.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
CP_READ_LATENCY or CPRDLAT

CP Reads attribute

Description
The number of read operations that are performed per second on this aggregate during a consistency point.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
CP_READS or CP_RD

File Rate attribute

Description
The number of file-based read and write operations that are performed per second on this aggregate.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

FILE_RATE

Latency attribute

Description

The latency (in microseconds) of all the operations that are performed on this aggregate.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

LATENCY or LAT

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

Node Key attribute **This attribute is a key attribute.**

Description

The resource key of the associated cluster-node.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

NODE_KEY

OverCommitted Percent attribute

Description

Overcommitted size (in percentage) of the aggregate.

Type

real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

PERCENT_OVERCOMMIT or PROVRCOM

Raid Size attribute

Description

Size of the redundant array of independent disks (RAID).

Type

integer (32-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

RAID_SIZE

Raid Status attribute

Description

Status of the RAID.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

RAID_STATUS or RAIDSTATUS

Size Available Percent attribute

Description

Available size of the aggregate in percentage.

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

SIZE_AVAILABLE_PERCENT or SZAVLPRCNT

Size Used Percent attribute

Description

Used size of the aggregate in percentage.

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

SIZE_USED_PERCENT or SZUSDPRCNT

Snapshot Size Available attribute

Description

Available size for the snapshot in GB.

Type

real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

SNAPSHOT_SIZE_AVAILABLE or SNPSZAVL

Snapshot Size Used attribute

Description

Size (in GB) that is used for snapshot.

Type

real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

SNAPSHOT_SIZE_USED or SNPSZUSD

Throughput attribute

Description

The total data that is sent and received (in bytes per second) by an aggregate.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

THROUGHPUT or THRPT

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 Committed attribute
Description
Total committed space (in GB).
Type
real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
TOTAL_COMMITTED or TOT_COMMTD

Total Size attribute
Description
Aggregate total size in GB.
Type
real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
TOTAL_SIZE or TOT_SIZE

Total Transfers attribute
Description
The total number of transfers that are serviced per second by the aggregate.
Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
TOTAL_TRANSFERS or TOTL_XFER

User Read Blocks attribute
Description
The number of blocks that are read per second from this aggregate
Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
USER_READ_BLOCKS or USR_RD_BLK

User Read Latency attribute
Description
The average latency (in microseconds) per block for the read operations that are performed by the user.
Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
USERREADLATENCY or USRRDLAT

User Reads attribute
Description
The number of read operations that are performed per second on this aggregate by the user.
Type
real number (64-bit gauge) with one decimal places of precision with enumerated

values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
USER_READS or USR_RD

User Write Blocks attribute

Description
The number of blocks that are written per second to this aggregate.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
USER_WRITE_BLOCKS or USR_WR_BLK

User Write Latency attribute

Description
The average latency (in microseconds) per block for the write operations that are performed by the user.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
USERWRITELATENCY or USRWTLAT

User Writes attribute

Description
The number of write operations that are performed per second on this aggregate by the user.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
USER_WRITES or USR_WR

Utilization attribute

Description
The CPU usage (in percentage) of this aggregate.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
UTILIZATION or UTIL

Cluster Node attribute group

The Cluster Node attribute group contains information about cluster nodes.

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 Cluster Node attribute group:

Average Processor Busy attribute

Description
The average processor usage (in percentage) across all processors in the system.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated

values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

AVERAGE_PROCESSOR_BUSY or AVG_PR_BSY

Cluster Key attribute **This attribute is a key attribute.**

Description

The resource key of the associated cluster.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CLUSTER_KEY or CLST_KEY

Cluster Node Name attribute **This attribute is a key attribute.**

Description

Name of the cluster node.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CLUSTER_NODE_NAME or CN_NAME

CPU Busy attribute

Description

The time (in percentage) for which one or more processors is busy in the system.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CPU_BUSY

Failed Fan Count attribute **This attribute is a key attribute.**

Description

The number of failed fans.

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

FAILED_FAN_COUNT or FLDFANCNT

Failed Power Supply Count attribute **This attribute is a key attribute.**

Description

The number of failed power supply units. Following are the possible values: "normal", "one_failed", "many_failed".

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

FAILED_POWER_SUPPLY_COUNT or FLDPWCNT

FailOver State attribute

Description

Storage failover configuration state.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

FAILOVER_STATE or FAILOVERST

Is Node Healthy attribute

Description

Indicates whether the cluster node is healthy.

Type

string with enumerated values. The following values are defined: unavailable (unavailable), Yes (true), No (false). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

IS_NODE_HEALTHY or CN_HEALTHY

Max Aggr Utilization attribute

Description

The maximum aggregate usage by this cluster node.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

MAX_AGGREGATE_UTILIZATION or MAX_AGG_UZ

Memory Size attribute **This attribute is a key attribute.**

Description

Memory size of node in GB.

Type

real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

MEMORY_SIZE or MEMORY_SZ

Net Data Recv attribute **This attribute is a key attribute.**

Description

The data that is received (in KB per second) over the network.

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

NET_DATA_RECV or CN_DT_RCV

Net Data Sent attribute **This attribute is a key attribute.**

Description

The data that is sent (in KB per second) over the network.

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

NET_DATA_SENT or CN_DT_SNT

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

Node Key attribute **This attribute is a key attribute.**
Description
The resource key of this cluster-node.
Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
NODE_KEY

Number Of Processors attribute
Description
Number of allocated processors.
Type
integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
NUMBER_OF_PROCESSORS or NUMOFPRCRS

NVRAM Battery Status attribute **This attribute is a key attribute.**
Description
Status of the NVRAM battery.
Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
NVRAM_BATTERY_STATUS or CN_NVR

Other Operations attribute
Description
The number of other fibre channel protocol (FCP) operations at the node level.
Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
OTHER_OPS or OTHR_OPS

Read Operations attribute
Description
The number of aggregated FCP read operations at the node level.
Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
READ_OPS or RD_OPS

Read Throughput attribute
Description
The data that is read (in KB per second) from this cluster node.
Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
READ_THROUGHPUT or RD_THRGPUT

Serial Number attribute

Description
Serial number of the node.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
SERIAL_NUMBER or SERIALNUM

Sys Avg Latency attribute

Description
The average latency (in microseconds) of the system.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
SYS_AVG_LATENCY or SYS_AVG_LT

Sys Read Latency attribute

Description
The read latency (in microseconds) of the system.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
SYS_READ_LATENCY or SYS_RD_LT

Sys Write Latency attribute

Description
The write latency (in microseconds) of the system.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
SYS_WRITE_LATENCY or SYS_WR_LT

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 Throughput attribute

Description
The total data that is sent and received (in KB per second) by the cluster node.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
TOTAL_THROUGHPUT or TO_THRGPUT

Utilization attribute

Description

The processor usage by this cluster node.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

UTILIZATION or UTLZN

Write Throughput attribute

Description

The data that is written (in KB per second) to this cluster node.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

WRITE_THROUGHPUT or WR_THRGPUT

Clusters attribute group

The Clusters attribute group contains information about clusters.

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 Clusters attribute group:

Cluster Key attribute This attribute is a key attribute.

Description

The resource key for this cluster.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CLUSTER_KEY or CLST_KEY

Cluster Name attribute This attribute is a key attribute.

Description

Name of the cluster.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CLUSTER_NAME or CLST_NAME

Is Metro Cluster attribute This attribute is a key attribute.

Description

Indicates whether this is a metro cluster.

Type

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

Warehouse name

IS_METRO_CLUSTER or METRO_CLST

Location attribute This attribute is a key attribute.

Description

The physical location of the cluster.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
LOCATION

Management IP attribute **This attribute is a key attribute.**

Description
Primary IP address of the cluster.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
MANAGEMENT_IP or MNGMNTIP

Max Aggregate Utilization attribute

Description
The maximum aggregate usage by this cluster in percentage.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
MAX_AGGREGATE_UTILIZATION or MAX_AGG_UZ

Max Node Utilization attribute

Description
The maximum node usage by this cluster in percentage.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
MAX_NODE_UTILIZATION or MAX_NOD_UZ

Metro Cluster Mode attribute **This attribute is a key attribute.**

Description
Mode of the metro cluster.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
METRO_CLUSTER_MODE or CLST_MODE

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

Status attribute

Description
Current status of the host that is based on all events.

Type
string with enumerated values. The following values are defined: not mapped

(not_mapped), ok (ok), ok with suppressed (ok_with_suppressed), degraded (degraded), unreachable (unreachable), unknown (unknown). Any value that does not have a definition here is displayed in the User Interface

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

Total Operations attribute

Description
The number of operations that are serviced per second by the cluster.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
TOTAL_OPS or TOTLOPS

Total Throughput attribute

Description
The total data that is sent and received (in KB per second) by the cluster.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
TOTAL_THROUGHPUT or TOTLTHRPT

Version attribute This attribute is a key attribute.

Description
Software version of the requested cluster.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
VERSION

DataSource attribute group

The Datasource attribute group contains information about the configured OnCommand API Services.

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 DataSource attribute group:

Agent Connection attribute

Description
Current status of the agent connection to the configured data source.

Type
string with enumerated values. The following values are defined: Down (false), Up (true). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
AGENT_CONNECTION or AGNCN

Data Source attribute **This attribute is a key attribute.**

Description
Name of the product that is being monitored.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
DATA_SOURCE or DS

IP / Host Name attribute **This attribute is a key attribute.**

Description
IP address of the OnCommand API Services.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
IP_ADDRESS or IP

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

Port attribute **This attribute is a key attribute.**

Description
Port number on which the OnCommand API Service runs.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
PORT

Refresh Interval attribute **This attribute is a key attribute.**

Description
The time interval (in minutes) after which the data provider sends a request to fetch data from the OnCommand API Services.

Type
string with enumerated values. The following values are defined: -1 (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
REFRESH_INTERVAL or RFINTRVL

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

Version attribute **This attribute is a key attribute.**

Description

Versions of the OnCommand API Services.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

VERSION

Disks attribute group

The Disks attribute group contains information about disks.

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 Disks attribute group:

Block Rate attribute

Description

The number of block read and write operations that occurred per second on the disk.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

BLOCK_RATE

Cluster Key attribute **This attribute is a key attribute.**

Description

The resource key of the associated cluster.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CLUSTER_KEY or CLST_KEY

Composite Latency attribute

Description

The sum of the read, write, and consistency point latency for the disk.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

COMPOSITE_LATENCY or CMP_LATNCY

Container Type attribute

Description

Type of the overlying disk container.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CONTAINER_TYPE or CONT_TYPE

CP Read Blocks attribute

Description

The number of blocks that are transferred per second to the disk for the consistency point read operations.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CP_READ_BLOCKS or CP_RD_BLK

CP Read Latency attribute

Description

The average latency (in milliseconds) per block for consistency point read operations that occurred from the disk.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CP_READ_LATENCY or CP_LATNCY

CP Reads attribute

Description

The number of read operations that occurred per second during the consistency point on the disk.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CP_READS or CP_RD

Disk Busy attribute

Description

The time in percentage since the last outstanding request to the disk.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

DISK_BUSY

Disk Key attribute **This attribute is a key attribute.**

Description

The resource key for this disk.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

DISK_KEY

Disk Name attribute **This attribute is a key attribute.**

Description

Name of the disk. The name will be displayed in the following format: "data disk 0b.18", "parity disk 0b.17", "dparity disk 0b.16", and so on.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
DISK_NAME

Failed Reason attribute

Description
Cuase of the disk failure.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
FAILED_REASON or FLD_REASON

File Rate attribute

Description
The number of file-based read and write operations that occurred per second on the disk.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
FILE_RATE

Guaranteed Read Latency attribute

Description
The average latency (in microseconds) per block for guaranteed read operations.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
GUARANTEED_READ_LATENCY or GRDLAT

Guaranteed Write Latency attribute

Description
The average latency (in microseconds) per block for guaranteed write operations.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
GUARANTEED_WRITE_LATENCY or GWTLAT

Is Offline attribute

Description
Indicates whether the disk is offline.

Type
string with enumerated values. The following values are defined: unavailable (unavailable), Yes (true), No (false). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
IS_OFFLINE or OFFLINE

Is Virtual attribute

Description
Indicates whether the disk is a virtual disk.

Type
string with enumerated values. The following values are defined: unavailable (unavailable), Yes (true), No (false). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
IS_VIRTUAL or VIRTUAL

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

Serial Number attribute

Description
Serial number of the disk.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
SERIAL_NUMBER or SERL_NUM

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 Bytes attribute

Description
Total size (in GB) of the disk.

Type
real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
TOTAL_BYTES or TOT_BYTES

Total Latency attribute

Description
The total latency (in microseconds) of the disk.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
TOTAL_LATENCY or TOTLAT

Total Operations attribute

Description
The number of operations that are serviced per second by the disk.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
TOTAL_OPS or TOTLOPS

Total Throughput attribute

Description

The total data that is sent and received (in bytes per second) from the disk.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

TOTAL_THROUGHPUT or TOTLTHRPT

Total Transfers attribute

Description

The total number of transfers that are serviced per second by the disk.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

TOTAL_TRANSFERS or TOTL_XFER

Used Bytes attribute

Description

Used size (in GB) of the disk.

Type

real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

USED_BYTES or USD_BYTES

User Read Blocks attribute

Description

The number of blocks that are transferred per second from the disk for the read operations that are performed by the user.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

USER_READ_BLOCKS or USR_RD_BLK

User Read Latency attribute

Description

The average latency (in milliseconds) per block for the read operations that are performed from a disk by the users.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

USER_READ_LATENCY or RD_LATNCY

User Reads attribute

Description

The number of read operations that are performed per second from the disk by the user.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
USER_READS or USR_RD

User Write Blocks attribute

Description
The number of blocks that are transferred per second from the disk for the write operations that are performed by the user.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
USER_WRITE_BLOCKS or USR_WR_BLK

User Write Latency attribute

Description
The average latency (in milliseconds) per block for the write operations that are performed to the disk by the users.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
USER_WRITE_LATENCY or WR_LATNCY

User Writes attribute

Description
The number of write operations that are performed per second on the disk by the user.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
USER_WRITES or USR_WR

Events attribute group

The Events attribute group contains information about events.

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 Events attribute group:

Event Key attribute This attribute is a key attribute.

Description
The resource key of the event.

Type
string with enumerated values. The following values are defined: Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
EVENT_KEY or EVTKEY

Event Severity attribute

Description
The severity of the event.

Type
string with enumerated values. The following values are defined: Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
EVENT_SEVERITY or EVTSEV

Event Source Key attribute

Description
The identifier of the resource on which the event is raised.

Type
string with enumerated values. The following values are defined: Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
EVENT_SOURCE_KEY or EVTSRCKEY

Event Status attribute **This attribute is a key attribute.**

Description
The status of the event.

Type
string with enumerated values. The following values are defined: Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
EVENT_STATUS or EVTSTAT

Event Time attribute **This attribute is a key attribute.**

Description
The time when the event occurred.

Type
string with enumerated values. The following values are defined: Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
EVENT_TIME or EVTTM

Event Type attribute

Description
The type or the class to which the event belongs.

Type
string with enumerated values. The following values are defined: Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
EVENT_TYPE or EVTTYPE

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

EventsOverview attribute group

The attribute group gives brief information on the count of events occurred per entity across the 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 EventsOverview attribute group:

Critical attribute **This attribute is a key attribute.**

Description

The number of events of the Critical type that occurred on the entity.

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CRITICAL or CR

Entity attribute **This attribute is a key attribute.**

Description

Name of the entity or object where event has triggered.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

ENTITY or EN

Error attribute **This attribute is a key attribute.**

Description

The number of events of the Error type that occurred on the entity.

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

ERROR or ER

Information attribute **This attribute is a key attribute.**

Description

The number of events of the Information type that occurred on the entity.

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

INFORMATION or IR

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

Normal attribute

Description
The number of events of the Error type that occurred on the entity.
Type
integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
NORMAL or NR

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

Warning attribute **This attribute is a key attribute.**
Description
The number of events of the Warning type that occurred on the entity.
Type
integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
WARNING or WR

LUNs attribute group

The LUNs attribute group contains information about LUNs.

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 LUNs attribute group:

Average Latency attribute
Description
The average time (in milliseconds) that is required to perform all the operations on this LUN.
Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
AVG_LATENCY or AVG_LATNCY

Is Mapped attribute
Description
Indicates whether the LUN is mapped.
Type
string with enumerated values. The following values are defined: unavailable (unavailable), Yes (true), No (false). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
IS_MAPPED or MAPPED

Is Online attribute
Description
Indicates whether the LUN is online.

Type

string with enumerated values. The following values are defined: unavailable (unavailable), Yes (true), No (false). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

IS_ONLINE or ONLINE

LUN Key attribute

Description

The resource key of this Lun.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

LUN_KEY

LUN Path attribute **This attribute is a key attribute.**

Description

Path name of the LUN including the volume or the Qtree where the LUN exists. The path name is similar to the following format: myvol/mylun or myvol/myqtree/mylun.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

LUN_PATH

LUN Size attribute

Description

LUN size in GB.

Type

real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

LUN_SIZE

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

Other Ops attribute

Description

The number of other operations that are performed per second on this LUN.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

OTHER_OPS

Qtree Key attribute

Description

The resource key of the associated Qtree.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

QTREE_KEY

Read Data attribute

Description

The data that is read (in KB per second) from this LUN.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

READ_DATA or RD_DATA

Read Ops attribute

Description

The number of read operations that are performed per second on this LUN.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

READ_OPS

Serial Number attribute

Description

Serial number of the LUN. The serial number is a 12-character string that consists of upper and lower-case letters, numbers, slashes (/), and hyphen (-) characters.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

SERIAL_NUMBER or SERL_NUM

Size Used attribute

Description

Size used in GB.

Type

real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

SIZE_USED

StorageVM Key attribute

Description

The resource key of the associated storage-vm

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

STORAGEVM_KEY or STRGVM_KEY

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 Ops attribute
Description
The number of operations that are performed per second by this LUN.
Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
TOTAL_OPS or TOTL_OPS

Volume Key attribute
Description
The resource key of the associated volume.
Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
VOLUME_KEY or VOL_KEY

Write Data attribute
Description
The data that written (in KB per second) to this LUN.
Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
WRITE_DATA or WR_DATA

Write Ops attribute
Description
The number of write operations that are performed per second on this LUN.
Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface
Warehouse name
WRITE_OPS

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:

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), DELETED CONFIGURATION (72). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

ERROR_CODE or ERRCODE

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

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

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

Qtrees attribute group

The Qtrees attribute group contains information about qtrees.

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 Qtrees attribute group:

Disk Limit attribute
Description
Disk limit for the Qtree.

Type
integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
DISK_LIMIT or DISK_LMT

Disk Used attribute

Description
Disk space (in GB) used.

Type
real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
DISK_USED

File Limit attribute

Description
File limit for the Qtree.

Type
integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
FILE_LIMIT or FILE_LMT

File Used attribute

Description
Files used (in GB) in Qtree.

Type
real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
FILE_USED

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

Qtree Key attribute

Description
The resource key of this Qtree.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
QTREE_KEY

Qtree Name attribute **This attribute is a key attribute.**

Description
Name of the Qtree.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
QTREE_NAME or QTR_NAME

Security Style attribute

Description
Security style of the Qtree

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
SECURITY_STYLE or SEC_STYLE

Soft File Limit attribute

Description
Displays the soft limit for the number of files that are permitted on the Qtree.

Type
integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
SOFT_FILE_LIMIT or SOFTFL_LMT

Status attribute

Description
Status of the Qtree.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
STATUS

StorageVM Key attribute

Description
The resource key of the associates Storage-VM.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
STORAGEVM_KEY or STRGVM_KEY

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

Volume Key attribute

Description
The resource key of the associates Volume.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
VOLUME_KEY or VOL_KEY

Volumes attribute group

The Volumes attribute group contains information about volumes.

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:

Aggregate Key attribute

Description
The resource key of the associated aggregate.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
AGGREGATE_KEY or AGGR_KEY

Average Latency attribute

Description
The average time (in microseconds) to perform all the operations on the volume.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
AVERAGE_LATENCY or AVGLAT

CIFS other latency attribute

Description
The average time (in microseconds) for the write anywhere file layout (WAFL) file system to process other common internet file system (CIFS) operations to this volume.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
CIFS_OTHER_LATENCY or CIFSOTLAT

CIFS Other Operations attribute

Description
The number of other CIFS operations that are performed on this volume.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
CIFS_OTHER_OPS or CIPSOPS

CIFS Read latency attribute

Description
The average time that is required for the WAFL file system to process CIFS read requests to the volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CIFS_READ_LATENCY or CIFSRDLAT

CIFS Read Operations attribute

Description

The number of CIFS read operations that are performed on this volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CIFS_READ_OPS or CIFSRDOPS

CIFS Write latency attribute

Description

The average time (in microseconds) for the CIFS write operations that are performed on this volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CIFS_WRITE_LATENCY or CIFSWTLAT

CIFS Write Operations attribute

Description

The number of CIFS write operations that are performed on the volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

CIFS_WRITE_OPS or CIFSWTOPS

iSCSI Read Operations attribute

Description

The total number of read operations for all the LUNs in the system that are accessed over iSCSI.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

ISCSI_READ_OPS or ISCSIROPS

NFS Other Latency attribute

Description

The average time (in microseconds) that is required to perform other network file system (NFS) operations on this volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

NFS_OTHER_LATENCY or NFSOLAT

NFS Other Operations attribute

Description
The number of other NFS operations that are performed on this volume.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
NFS_OTHER_OPS or NFSOTOPS

NFS Read latency attribute

Description
The average time (in microseconds) for the NFS read operations that are performed on this volume.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
NFS_READ_LATENCY or NFSRDLAT

NFS Read Operations attribute

Description
The number of NFS read operations that are performed on this volume.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
NFS_READ_OPS or NFSRDOPS

NFS Write Latency attribute

Description
The average time (in microseconds) that is required to perform NFS write operations on this volume.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
NFS_WRITE_LATENCY or NFSRTOPS

NFS Write Operations attribute

Description
The number of NFS write operations that are performed per second on this volume.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
NFS_WRITE_OPS or NFSROPS

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

Other Latency attribute

Description

The average time (in milliseconds) that is required to perform other operations on this volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

OTHER_LATENCY or OTH_LATNCY

Other Ops attribute

Description

The number of other operations that are performed per second on this volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

OTHER_OPS or OTH_OPS

Overwrite Reserve Actual Used attribute

Description

Size used (in MB) for overwrite reserved space.

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

OVERWRITE_RESERVE_ACTUAL_USED or OVR_ACUSED

Quota Over Committed attribute

Description

Displays the amount of physical space in GB in the Qtrees that can be used before the system generates a Volume Quota Overcommitted event.

Type

real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

QUOTA_OVER_COMMITTED or QTAOVRCMTD

Quota Status attribute

Description

Status of the quota.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

QUOTA_STATUS or QUOTA_STS

Read Latency attribute

Description

The average time (in milliseconds) that is required to perform read operations on this volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
READ_LATENCY or RD_LATNCY

Read Ops attribute

Description
The number of read operations that are performed per second on this volume.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
READ_OPS or RD_OPS

Read Throughput attribute

Description
The total data that is received (in bytes per second) by this volume.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
READ_THROUGHPUT or RDTTHRPT

SIS Status attribute

Description
Indicates the status and progress of the deduplication operation on a FlexVol volume.

Type
string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
SIS_STATUS

Size Total attribute

Description
Total size in GB in the active file system (total volume minus the snapshot reserve).

Type
real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
SIZE_TOTAL or SIZE_TOT

Size Used attribute

Description
Size used in GB to hold active file system data. This value is reported by "df" as used for the volume. The value includes data, hole reserves, overwrite reserves and snapshot overflow.

Type
real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
SIZE_USED

Size Avail Percent attribute

Description
Percentage of bytes that are available to hold active file system data.

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

SIZE_AVAIL_PERCENT or SZ_AVL_PER

Size Used Percent attribute

Description

Percentage of bytes that are used to hold active file system data.

Type

integer (64-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

SIZE_USED_PERCENT or SZ_USD_PER

Size Used By Snapshots attribute

Description

Reports size that is used (in GB) by snapshots.

Type

real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

SIZE_USED_BY_SANPSHOTS or SZUSDBYSNP

Snapshot Reserve Percent attribute

Description

Percentage of snapshot reserve.

Type

integer (32-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

PERCENTAGE_SNAPSHOT_RESERVE or SNP_USED

Snapshot Reserve Percent Used attribute

Description

Percentage of bytes that are used to hold snapshot data. This value can be greater than the snapshot reserve size, but does not include the overwrite reserve space.

Type

integer (32-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

PERCENTAGE_SNAPSHOT_RESERVE_USED or SNP_PCT

StorageVM Key attribute

Description

The resource key of the associated storage volume.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

STORAGEVM_KEY or STRGVM_KEY

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 Inode Files attribute

Description

Number of inodes of this volume.

Type

integer (32-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

INODE_FILES_TOTAL or INDFILETOT

Total Ops attribute

Description

The number of operations that are performed per second by this volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

TOTAL_OPS or TOTL_OPS

Total throughput attribute

Description

The total data that is sent and received (in bytes per second) by this volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

TOTAL_THROUGHPUT or TOTTHRPT

Used Inode Files attribute

Description

Number of inodes that are used for this volume.

Type

integer (32-bit gauge) with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

INODE_FILES_USED or INDFILEUSD

Volume Key attribute

Description

The resource key of this Volume.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

VOLUME_KEY or VOL_KEY

Volume Name attribute **This attribute is a key attribute.**

Description

Name of the volume.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

VOLUME_NAME or VOL_NAME

Volume Size attribute

Description

Size of volume in GB including the snapshot reserve.

Type

real number (64-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

VOLUME_SIZE or VOL_SIZE

Volume State attribute

Description

Current state of the volume.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

VOLUME_STATE or VOL_STATE

Volume Style attribute

Description

Style of a volume. Following are the possible values: 'not_mapped' or ' flexvol' or ' striped' or ' infinivol.

Type

string with enumerated values. The following values are defined: unavailable (unavailable). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

VOLUME_STYLE or VOL_STYLE

Write Latency attribute

Description

The average time (in milliseconds) that is required to perform write operations on this volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

WRITE_LATENCY or WR_LATNCY

Write Ops attribute

Description

The number of write operations that are performed per second on this volume.

Type

real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name

WRITE_OPS or WR_OPS

Write Throughput attribute

Description

The total data that is sent (in bytes per second) by this volume.

Type
real number (64-bit gauge) with one decimal places of precision with enumerated values. The following values are defined: unavailable (-1). Any value that does not have a definition here is displayed in the User Interface

Warehouse name
WRITE_THROUGHPUT or WRTTHRPT

Disk capacity planning for historical data

Disk capacity planning for a monitoring agent is a prediction of the amount of disk space to be consumed for each attribute group with historical data that is being collected. Required disk storage is an important factor 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 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 19.

Attribute group
Name of the attribute group used to create the table in the warehouse database if it is short enough to fit in the table naming constraints of the database 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 19.

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

Database bytes per row (warehouse)
Estimate of the record length for detailed records written to the warehouse database, if the attribute group is configured to be written to the warehouse. Detailed records are 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 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 is monitoring each processor in your computer and you have a dual processor computer, the number of rows is two.

Table 1. Capacity planning for historical data logged by the NetApp Storage agent

Table	Attribute group	Bytes per row (agent)	Database bytes per row (warehouse)	Aggregate bytes per row (warehouse)
KNU02AGREG	KNU_AGGREGATES	376	511	1418
KNU05DISK	KNU_DISKS	550	688	1568
KNU04LUN	KNU_LUNS	376	464	1038
KNUPOBJST	KNU_PERFORMANCE_OBJECT_STATUS	288	289	326

Table 1. Capacity planning for historical data logged by the NetApp Storage agent (continued)

Table	Attribute group	Bytes per row (agent)	Database bytes per row (warehouse)	Aggregate bytes per row (warehouse)
KNU06QTREE	KNU_QTREES	372	431	938
KNU03VOL	KNU_VOLUMES	560	842	2580

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

Chapter 5. Situations reference

A situation is a logical expression involving one or more system conditions. Situations are used to monitor the condition of systems in your network. You can manage situations from the Tivoli Enterprise Portal by using the Situation Editor or from the command-line interface using the tacmd commands for situations. You can manage private situations in the private configuration XML file.

About situations

The monitoring agents that you use to monitor your system environment include a set of predefined situations that you can use as-is. You can also create new situations to meet your requirements.

Predefined situations contain attributes that check for system conditions common to many enterprises. Using predefined situations can improve the speed with which you can begin using the IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage. 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 NetApp Storage agent managed systems are assigned by default.

Expert advice
Comments and instructions to be read in the event workspace.

Action
Command to be sent to the system.

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

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

Additional information about situations

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

For a list of the predefined situations and information about each individual situation for this monitoring agent, see "Predefined situations."

Predefined situations

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

- NetApp Storage
  - Not applicable
- Aggregates
  - KNU_AggregateStorage_Critical

  - KNU_Aggregate_Offline
- Disks
  - Not applicable
- Events
  - Not applicable
- LUNs
  - Not applicable
- Monitored DataSource
  - KNU_Agent_Down
  - KNU_Critical_Alert
  - KNU_Error_Alert
  - KNU_Warning_Alert
- Qtrees
  - Not applicable
- Volumes
  - KNU_Volume_Offline

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.

Aggregates Navigator item

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

KNU_AggregateStorage_Critical situation

Description

Percentage of OverCommitted usage on this aggregate is above normal.

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

Formula

```
*IF *VALUE KNU_AGGREGATES.Percent_OverCommit *NE 'unavailable' *AND *VALUE
KNU_AGGREGATES.Percent_OverCommit *GT 89
```

See "Attributes in each attribute group" on page 20 for descriptions of the attributes in this formula.

Distribution

This situation is automatically distributed to instances of this agent.

Run at startup

Yes

Sampling interval

30 seconds

Situation persistence

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

Error conditions

Critical

Clearing conditions

The situation clears when the condition becomes false.

KNU_Aggregate_Offline situation

Description

Aggregate is offline.

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

Formula

```
*IF *VALUE KNU_AGGREGATES.AGGR_STATE *NE 'unavailable' *AND *VALUE
KNU_AGGREGATES.AGGR_STATE *EQ 'offline'
```

See "Attributes in each attribute group" on page 20 for descriptions of the attributes in this formula.

Distribution

This situation is automatically distributed to instances of this agent.

Run at startup

Yes

Sampling interval

30 seconds

Situation persistence

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

Error conditions

Critical

Clearing conditions

The situation clears when the condition becomes false.

Disks Navigator item

No predefined situations are included for this Navigator item.

Events Navigator item

No predefined situations are included for this Navigator item.

LUNs Navigator item

No predefined situations are included for this Navigator item.

Monitored DataSource Navigator item

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

KNU_Agent_Down situation

Description

The agent connection is down.

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

Formula

```
*IF *VALUE KNU_DATASOURCE.Agent_Connection *EQ 'Down'
```

See "Attributes in each attribute group" on page 20 for descriptions of the attributes in this formula.

Distribution

This situation is automatically distributed to instances of this agent.

Run at startup

Yes

Sampling interval

30 seconds

Situation persistence

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

Error conditions

Critical

Clearing conditions

The situation clears when the condition becomes false.

KNU_Critical_Alert situation

Description

Critical Alerts in the environment.

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

Formula

```
*IF *VALUE KNU_EVENTSOVERVIEW.Critical *GE 1
```

See "Attributes in each attribute group" on page 20 for descriptions of the attributes in this formula.

Distribution

This situation is automatically distributed to instances of this agent.

Run at startup

Yes

Sampling interval

30 seconds

Situation persistence

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

Error conditions

Critical

Clearing conditions

The situation clears when the condition becomes false.

KNU_Error_Alert situation

Description

Error Alerts in the environment.

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

Formula

```
*IF *VALUE KNU_EVENTSOVERVIEW.Error *GE 1
```

See "Attributes in each attribute group" on page 20 for descriptions of the attributes in this formula.

Distribution

This situation is automatically distributed to instances of this agent.

Run at startup

Yes

Sampling interval

30 seconds

Situation persistence

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

Error conditions

Fatal

Clearing conditions

The situation clears when the condition becomes false.

KNU_Warning_Alert situation

Description

Warning Alerts in the environment.

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

Formula

```
*IF *VALUE KNU_EVENTSOVERVIEW.Warning *GE 1
```

See "Attributes in each attribute group" on page 20 for descriptions of the attributes in this formula.

Distribution

This situation is automatically distributed to instances of this agent.

Run at startup

Yes

Sampling interval

30 seconds

Situation persistence

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

Error conditions

Warning

Clearing conditions

The situation clears when the condition becomes false.

Qtrees Navigator item

No predefined situations are included for this Navigator item.

Volumes Navigator item

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

KNU_Volume_Offline situation

Description

Volume is offline.

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

Formula

```
*IF *VALUE KNU_VOLUMES.Volume_State *NE 'unavailable' *AND *VALUE
KNU_VOLUMES.Volume_State *EQ 'offline'
```

See “Attributes in each attribute group” on page 20 for descriptions of the attributes in this formula.

Distribution
This situation is automatically distributed to instances of this agent.

Run at startup
Yes

Sampling interval
30 seconds

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

Error conditions
Critical

Clearing conditions
The situation clears when the condition becomes false.

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 the *Tivoli Enterprise Portal User's Guide*.

Predefined Take Action commands

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

The IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage 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.

About policies

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.

Additional information about policies

This monitoring agent does not provide predefined policies. For more information about working with policies, see 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.

Chapter 8. Troubleshooting

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

See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information. Also see "Support information" on page 94 for other problem-solving options.

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.

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 72
- "Consulting the lists of identified problems and workarounds" on page 72

Gathering product information for IBM Software Support

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

Table 2. Information to gather before contacting IBM Software Support

Information type	Description
Log files	Collect trace log files from failing systems. Most logs are located in a logs subdirectory on the host computer. See "Principal trace log files" on page 73 for lists of all trace log files and their locations. See the *Tivoli Enterprise Portal User's Guide* for general information about the IBM Tivoli Monitoring environment.
Operating system	Operating system version number and patch level
Messages	Messages and other information displayed on the screen
Version numbers for IBM Tivoli Monitoring	Version number of the following members of the monitoring environment: • IBM Tivoli Monitoring. Also provide the patch level, if available. • IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage
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 *install_dir*/bin directory, where *install_dir* 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. See the "pdcollect tool" section in the "Tools" chapter of the *IBM Tivoli Monitoring Troubleshooting Guide* for more information about using this tool.

See http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_(general) for information about working with IBM Software Support.

Using logging

Logging is the primary troubleshooting feature in the NetApp Storage agent. *Logging* refers to the text messages and trace data that is generated by the NetApp Storage agent. Messages and trace data are sent to a file.

Trace data captures transient information about the current operating environment when a component or application fails to operate as designed. IBM Software Support personnel use the captured trace information to determine the source of an error or unexpected condition. See "Trace logging" for more information.

Consulting the lists of identified problems and workarounds

Known problems have been 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 83.

See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

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 73
- "Examples: Using trace logs" on page 76
- "Setting RAS trace parameters by using the GUI" on page 78

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

Log files have naming conventions.

Log file naming conventions

Table 3 on page 73 provides the names, locations, and descriptions of RAS1 log files. The log file names adhere to the following naming convention:

Windows systems

hostname_productcode_program_HEXtimestamp-nn.log

Linux and UNIX systems

hostname_productcode_HEXtimestamp-nn.log

where:

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

productcode
Two-character product code. For IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage, the product code is nu.

program
Name of the program being run.

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

nn Rolling log suffix.

Principal trace log files

Trace log files are located on various systems.

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

Table 3. Trace log files for troubleshooting agents

System where log is located	File name and path	Description
On the Tivoli Enterprise Monitoring Server	• **Windows:** The file in the *install_dir*\InstallITM path • **UNIX:** The candle_installation.log file in the *install_dir*/logs path • **Linux:** The candle_installation.log file in the *install_dir*/logs path	Provides details about products that are installed. **Note:** Trace logging is enabled by default. A configuration step is not required to enable this tracing.
On the Tivoli Enterprise Monitoring Server	The Warehouse_Configuration.log file is in the following location on Windows systems: *install_dir*\InstallITM	Provides details about the configuration of data warehousing for historical reporting.

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

System where log is located	File name and path	Description
On the Tivoli Enterprise Monitoring Server	The name of the RAS log file is as follows: • **Windows:** *install_dir*\logs\ *hostname*_ms_*timestamp-nn*.log • **UNIX:** *install_dir*/logs/ *hostname*_ms_*timestamp-nn*.log • **Linux:** *install_dir*/logs/ *hostname*_ms_*timestamp-nn*.log **Note:** File names for RAS1 logs include a hexadecimal time stamp. Also on UNIX systems, a log with a decimal time stamp is provided: *hostname_productcode_timestamp*.log and *hostname_productcode_ timestamp*.pid nnnnn in the *install_dir*/logs path, where *nnnnn* is the process ID number.	Traces activity on the monitoring server.
On the Tivoli Enterprise Portal Server	The name of the RAS log file is as follows: • **Windows:** *install_dir*\logs\ *hostname* _cq_HEXtimestamp-nn.log • **UNIX:** *install_dir* /logs/*hostname*_cq_HEXtimestamp-nn.log • **Linux:** *install_dir* /logs/*hostname*_cq_HEXtimestamp-nn.log **Note:** File names for RAS1 logs include a hexadecimal time stamp. Also on UNIX systems, a log with a decimal time stamp is provided: *hostname_productcode_timestamp* .logand *hostname_productcode_ timestamp*.pid*nnnnn* in the *install_dir*/logs path, where *nnnnn* is the process ID number.	Traces activity on the portal server.
On the Tivoli Enterprise Portal Server	The teps_odbc.log file is located in the following path: • **Windows:** *install_dir*\InstallITM • **UNIX:** *install_dir*/logs • **Linux:** *install_dir*/logs	When you enable historical reporting, this log file traces the status of the warehouse proxy agent.

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

System where log is located	File name and path	Description
On the computer that hosts the monitoring agent	The RAS1 log files are as follows: • **Windows:** *hostname*_nu_*instance_name*_knuagent_HEXtimestamp-nn.log in the *install_dir*\tmaitm6\logs directory • **UNIX:** *hostname*_nu_*instance_name*_knuagent_HEXtimestamp-nn.log in the *install_dir*/logs directory • **Linux:** *hostname*_nu_*instance_name*_knuagent_HEXtimestamp-nn.log in the *install_dir*/logs directory These logs are in the following directories: • **Windows:** *install_dir*\tmaitm6\logs • **UNIX:** *install_dir*/logs • **Linux:** *install_dir*/logs On Linux systems, the following additional logs are provided: – *hostname*_nu_timestamp.log – *hostname*_nu_timestamp.pidnnnnn in the *install_dir*/logs path, where *nnnnn* is the process ID number	Traces activity of the monitoring agent.
On the computer that hosts the monitoring agent	The agent operations log files are as follows: *instance_hostname*NU.LG0 is the current log created when the agent was started. *instance_hostname*_NU.LG1 is the backup of the previous log. These logs are in the following directory depending on the operating system that you are using: • **Windows:** *install_dir*\tmaitm6\logs • **Linux:** *install_dir*/logs • **UNIX:** *install_dir*/logs	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. 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 *previous* monitoring session: • Status of connectivity with the monitoring server • Situations that were running • The success or failure status of Take Action commands

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

System where log is located	File name and path	Description
On the computer that hosts the monitoring agent	The Take Action command log files are as follows: • *host*_nu_*instance*_*takeactioncommand* .log The logs are in the following directories: • **Windows:** *install_dir*\tmaitm6\logs • **UNIX:** *install_dir* /logs • **Linux:** *install_dir* /logs	Traces activity each time a Take Action command runs. For example, when a hypothetical **start_command** Take Action command runs, IBM Tivoli Monitoring generates a start_command.log file.
On the computer that hosts the monitoring agent	The data provider log files are as follows: • knu_data_provider_instance_startup.log • knu_data_provider_instance_n.log The logs are in the following directories: • **Windows:** *install_dir*\tmaitm6\logs • **Linux:** *install_dir*/logs	Traces monitoring agent data provider state and operations.
Definitions of variables: • *timestamp* is a time stamp with a format that includes year (y), month (m), day (d), hour (h), and minute (m), as follows: **yyyymmdd hhmm** • *HEXtimestamp* is a hexadecimal representation of the time at which the process was started. • *install_dir* represents the directory path where you installed the IBM Tivoli Monitoring component. *install_dir* can represent a path on the computer that hosts the monitoring system, the monitoring agent, or the portal. • *instance* refers to the name of the database instance that you are monitoring. • *instance_name* refers to the name of the agent instance. • *hostname* refers to the name of the computer on which the IBM Tivoli Monitoring component runs. • *nn* 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. • *productcode* specifies the product code, for example, um for Universal Agent or nt for Windows systems.		

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

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

Example two

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

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

See the following key points about the preceding excerpts:

- The monitoring server appends the **NU** product code to the server name to form a unique name (SERVER5B:NU) for this instance of the IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage. 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” on page 78 provide these entries.

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

1. In the Windows **Start** menu, click **Program Files** > **IBM Tivoli Monitoring** > **Manage Tivoli Enterprise Monitoring Services**. The Manage Tivoli Enterprise Monitoring Services window is displayed.
2. Right-click a component and click **Advanced** > **View Trace Log** in the menu. For example, if you want to view the trace log of the IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage, 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

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

See “Overview of log file management” on page 72 to ensure that you understand log rolling and can reference the correct log files when you manage log file generation.

In the log monitoring component of the PeopleSoft Application Domain agent and PeopleSoft Process Scheduler agent, a file name-based scheme is used to determine the newest PeopleSoft log file to monitor. For example, the `APPSRV_1016 log` file from October 16 appears to be newer than the `APPSRV_0218` file from February 18; thus, the `APPSRV_1016` file is monitored. This scheme becomes a problem when the date rolls over from December to January upon a new year. The log monitoring incorrectly identifies

APPSRV_1231 as a newer log file than APPSRV_0101. The suggested solution is to archive log files from the previous year into a different location. Keep only files from the current year in the directories containing the PeopleSoft domain and scheduler log files.

Setting RAS trace parameters by using the GUI

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

About this task

The IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage uses RAS1 tracing and generates the logs described in Table 3 on page 73. The default RAS1 trace level is ERROR. The default RAS1 trace level is ERROR.

Procedure

1. Open the Manage Tivoli Enterprise Monitoring Services window.
2. Select **Advanced** > **Edit Trace Parms**. 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:kqz ALL)
   - Maximum error tracing. KBB_RAS1=ERROR (UNIT:kqz 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_DEBUGd Setting** menu to log information that can help you diagnose communications and connectivity problems between the monitoring agent and the monitoring server. The **KDC_DEBUG** setting and the **Maximum error tracing** setting can generate a large amount of trace logging. Use these settings only temporarily, while you are troubleshooting problems. Otherwise, the logs can occupy excessive amounts of hard disk space.
8. Click **OK**. You see a message reporting a restart of the monitoring agent so that your changes take effect.

What to do next

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

Regularly prune log files other than the RAS1 log files in the logs directory. Unlike the RAS1 log files that are pruned automatically, other log types can grow indefinitely, for example, the logs in Table 3 on page 73 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 IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage uses RAS1 tracing and generates the logs described in Table 3 on page 73. The default RAS1 trace level is ERROR. The default RAS1 trace level is ERROR.

Procedure

1. Open the trace options file.
   - Windows:

     install_dir\tmaitm6\KNUENV
   - *install_dir* /config/nu.ini
2. Edit the line that begins with KBB_RAS1= to set trace logging preferences. For example, if you want detailed trace logging, set the **Maximum Tracing** option: KBB_RAS1=ERROR (UNIT:kqz 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 given program. When this value is exceeded, the oldest log files are discarded. The default value is 9.
   - **LIMIT**: The maximum size, in megabytes (MB) of a RAS1 log file. The default value is 5.
   - IBM Software Support might guide you to modify the following parameters:
     - **COUNT**: The number of log files to keep in the rolling cycle of one program startup. The default is 3.
     - **PRESERVE**: The number of files that are not to be reused in the rolling cycle of one program startup. The default value is 1.

     Note: The **KBB_RAS1_LOG** parameter also provides for the specification of the log file directory, log file name, and the inventory control file directory and name. Do not modify these values or log information can be lost.
4. Restart the monitoring agent so that your changes take effect.

What to do next

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

Regularly prune log files other than the RAS1 log files in the logs directory. Unlike the RAS1 log files that are pruned automatically, other log types can grow indefinitely, for example, the logs in Table 3 on page 73 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, with the exception of 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. It is the equivalent to setting `"Error Detail Flow State Input Output Metrics"`.

ANY
: Turns off tracing.

Detail
: Displays detailed information about each function.

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

ERROR
Logs internal error conditions.

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

Flow
Displays control flow data for each function entry and exit.

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

INPUT
Displays input data for each function.

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

Metrics
Displays metrics on each function.

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

OUTPUT
Displays output data for each function.

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

State
Displays the status for each function.

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

Example

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

```
kbbcre1.c, 400, May 29 2007, 12:54:43, 1.1, *
kbbcrn1.c, 400, May 29 2007, 12:54:42, 1.1, *
kdhb1de.c, 400, May 29 2007, 12:59:34, 1.1, KDH
kdh0med.c, 400, May 29 2007, 12:59:24, 1.1, KDH
kdhsrej.c, 400, May 29 2007, 13:00:06, 1.5, KDH
kdhb1fh.c, 400, May 29 2007, 12:59:33, 1.1, KDH
kdhb1oe.c, 400, May 29 2007, 12:59:38, 1.2, KDH
kdhs1ns.c, 400, May 29 2007, 13:00:08, 1.3, KDH
kbbacdl.c, 400, May 29 2007, 12:54:27, 1.2, ACF1
kbbaclc.c, 400, May 29 2007, 12:54:27, 1.4, ACF1
kbbac1i.c, 400, May 29 2007, 12:54:28, 1.11, ACF1
vkdhsfcn.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, it displays the information 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 ? to display a list of the supported commands. **ras1** is the command to modify trace settings. If you type it in the field at the end of the window, 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 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 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 kbbcrcd class of the Windows OS agent, the command is:
   ```
   ras1 set (UNIT:kbbcrcd 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 NetApp Storage agent, for example installation and configuration problems and workspace problems.

Note: You can resolve some problems by ensuring that your system matches the system requirements listed in the Prerequisites topic for the agent in the information center.

See the *IBM Tivoli Monitoring Troubleshooting Guide* for general troubleshooting information.

Installation and configuration troubleshooting

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

The problems and solutions in Table 4 on page 84 can occur during installation, configuration, and uninstallation of the agent.

Table 4. Problems and solutions for installation and configuration

Problem	Solution
(*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. **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.	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.
Diagnosing problems with product browse settings (Windows systems only).	When you have problems with browse settings, complete the following steps: 1. Click **Start** > **Programs** > **IBM Tivoli Monitoring** > **Manage Tivoli Enterprise Monitoring Services**. The Manage Tivoli Enterprise Monitoring Services window is displayed. 2. Right-click the Windows agent and select **Browse Settings**. A text window is displayed. 3. Click **Save As** and save the information in the text file. If requested, you can forward this file to IBM Software Support for analysis.
A message similar to `"Unable to find running CMS on CT_CMSLIST"` in the log file is displayed.	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: • 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.
The system is experiencing high CPU usage.	**Agent process:** View the memory usage of the KNUCMA process. If CPU usage seems to be excessive, restart the monitoring agent. **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.

Table 5. General problems and solutions for uninstallation

Problem	Solution
On Windows systems, uninstallation of IBM Tivoli Monitoring fails to uninstall the entire environment.	Be sure that you follow the general uninstallation process described in the *IBM Tivoli Monitoring Installation and Setup Guide*: 1. Remove Tivoli Enterprise Monitoring Server Application support by completing the following steps: a. Use Manage Tivoli Enterprise Monitoring Services. b. Select **Tivoli Enterprise Monitoring Server**. c. Right-click and select **Advanced**. d. Select **Remove TEMS application support**. e. Select the agent to remove its application support. 2. Uninstall the monitoring agents first, as in the following examples: • Uninstall a single monitoring agent for a specific database. -OR- • Uninstall all instances of a monitoring product, such as IBM Tivoli Monitoring for Databases. 3. Uninstall IBM Tivoli Monitoring.
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: 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**. To uninstall the monitoring agent, use the procedure described in the *IBM Tivoli Monitoring Installation and Setup Guide*.

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

Problem	Solution
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.	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 (*subsystem_name*:*hostname*:KNU). **Note:** When you monitor a multinode system, such as a database, IBM Tivoli Monitoring adds a subsystem name to the concatenated name, typically a database instance name. The length of the name that IBM Tivoli Monitoring generates is limited to 32 characters. Truncation can result in multiple components having the same 32-character name. If this problem happens, shorten the *hostname* portion of the name as follows: 1. Open the configuration file for the monitoring agent, which is located in the following path: • **On Windows:** *install_dir*\tmaitm6\K*product_code*CMA.INI. For example, the product code for the Monitoring Agent for Windows OS is NT. The file name is `KNTCMA.INI`. • **On UNIX and Linux:** `itm_home/config/product_code.ini` and `product_code.config`. For example, the file names for the Monitoring Agent for UNIX OS is `ux.ini` and `ux.config`. 2. Find the line 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 KNU, cannot be longer than 32 characters. **Note:** You must ensure that the resulting name is unique with respect to any existing monitoring component that was previously registered with the Tivoli Enterprise Monitoring Server. 4. Save the file. 5. Restart the agent.
The software inventory tag for the agent on UNIX and Linux systems is not removed during uninstallation of the agent.	After uninstalling the agent, manually remove the file named *`full name of agent`*`.cmptag` from the `$CANDLEHOME/properties/version/` directory.

Remote deployment troubleshooting

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

Table 6 on page 87 contains problems and solutions related to remote deployment.

Table 6. Remote deployment problems and solutions

Problem	Solution
While you are using the remote deployment feature to install the IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage, an empty command window is displayed on the target computer. This problem occurs when the target of remote deployment is a Windows computer. (See the *IBM Tivoli Monitoring Installation and Setup Guide* for more information about the remote deployment feature.)	Do not close or modify this window. It is part of the installation process and is dismissed automatically.
The removal of a monitoring agent fails when you use the remote removal process in the Tivoli Enterprise Portal desktop or browser.	This problem might occur when you attempt the remote removal process immediately after you have restarted the Tivoli Enterprise Monitoring Server. You must allow time for the monitoring agent to refresh its connection with the Tivoli Enterprise Monitoring Server before you begin the remote removal process.
Remote deployment of the agent to a 64-bit Windows server fails with a "time out" problem. Subsequent retries fail with a "file not transmitted" error.	The 64-bit Window OS agent was installed at the endpoint using the local installer instead of the IBM Tivoli Monitoring **createNode** command. If an **addSystem** command is used to deploy a 64-bit agent, the installation process loops continuously. This looping is caused by a perceived 32/64 bit compatibility (AC) component not being installed correctly. The install process running at the endpoint must be manually terminated. The remote deployment can now be executed by installing the AC component either locally or remotely. The agent can now be successfully deployed by running the **addSystem** command. If the agent installation is done locally, a pop-up menu is displayed indicating the 32/64 compatability component must be installed.

Agent troubleshooting

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

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

Table 7. Agent problems and solutions

Problem	Solution
Log data accumulates too rapidly.	Check the RAS trace option settings, which are described in "Setting RAS trace parameters by using the GUI" on page 78. The trace option settings that you can set on the `KBB_RAS1=` and `KDC_DEBUG=` lines potentially generate large amounts of data.
When using the **itmcmd agent** commands to start or stop this monitoring agent, you receive the following error message: `MKCIIN0201E Specified product is not configured.`	Include the command option **-o** to specify the instance to start or stop. The instance name must match the name used for configuring the agent. For example: `./itmcmd agent -o Test1 start nu` For more information about using the itmcmd commands, see the *IBM Tivoli Monitoring Command Reference*.

Table 7. Agent problems and solutions (continued)

Problem	Solution
A configured and running instance of the monitoring agent is not displayed in the Tivoli Enterprise Portal, but other instances of the monitoring agent on the same system are displayed in the portal.	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. IP.PIPE is the name given to Tivoli TCP/IP protocol for RPCs. The RPCs are socket-based operations that use TCP/IP ports to form socket addresses. IP.PIPE implements virtual sockets and multiplexes all virtual socket traffic across a single physical TCP/IP port (visible from the **netstat** command). A Tivoli process derives the physical port for IP.PIPE communications based on the configured, well-known port for the hub Tivoli Enterprise Monitoring Server. (This well-known port or BASE_PORT is configured by using the 'PORT:' keyword on the KDC_FAMILIES / KDE_TRANSPORT environment variable and defaults to '1918'.) The physical port allocation method is defined as (BASE_PORT + 4096*N), where N=0 for a Tivoli Enterprise Monitoring Server process and N={1, 2, ..., 15} for another type of monitoring server process. Two architectural limits result as a consequence of the physical port allocation method: • No more than one Tivoli Enterprise Monitoring Server reporting to a specific Tivoli Enterprise Monitoring Server hub can be active on a system image. • No more than 15 IP.PIPE processes can be active on a single system image. A single system image can support any number of Tivoli Enterprise Monitoring Server processes (address spaces) 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 simplified to one Tivoli Enterprise Monitoring Server per system image. No more than 15 IP.PIPE processes or address spaces can be active on a single system image. With the first limit expressed above, this second limitation refers specifically to Tivoli Enterprise Monitoring Agent processes: no more than 15 agents per system image. Continued on next row.

Table 7. Agent problems and solutions (continued)

Problem	Solution
Continued from previous row.	This limitation can be circumvented (at current maintenance levels, IBM Tivoli Monitoring V6.1, Fix Pack 4 and later) if the Tivoli Enterprise Monitoring Agent process is configured to use the EPHEMERAL IP.PIPE process. (This process is IP.PIPE configured with the 'EPHEMERAL:Y' keyword in the KDC_FAMILIES / KDE_TRANSPORT environment variable). The number of ephemeral IP.PIPE connections per system image has no limitation. If ephemeral endpoints are used, the Warehouse Proxy Agent is accessible from the Tivoli Enterprise Monitoring Server associated with the agents using ephemeral connections either by running the Warehouse Proxy Agent on the same computer or by using the Firewall Gateway feature. (The Firewall Gateway feature relays the Warehouse Proxy Agent connection from the Tivoli Enterprise Monitoring Server computer to the Warehouse Proxy Agent computer if the Warehouse Proxy Agent cannot coexist on the same computer.)
In the data provider logs there are outofmemory exception and if agent seems to be down then it indicates that data provider is struggling to get the essential resources of memory/CPU and allocation of heap size need to be increase for an instance.	To reconfigure the heap size allocation: 1. Stop the agent. 2. Edit the file. On Windows Systems 32-bit Agent system architecture: `TMAITM6` 64-bit Agent system architecture: `TMAITM6_x64` a. Go to the `CANDLE_HOME\<architecture>\KNUENV_INSTANCE_NAME` file. b. Search for the "ENVIRONMENT VARIABLES SPECIFIC TO THE AGENT" section and add the following variable: SET KNU_CUSTOM_JVM_ARGS=-Xmx512m On Linux Systems 32-bit Agent system architecture: `li6263` 64-bit Agent system architecture: `lx8266` a. Go to the `$CANDLEHOME/config/nu/bin/nu_INSTANCE_NAME.config` file. b. Search for the "Agent Specific Environment Variables" section and add the following variable: KNU_CUSTOM_JVM_ARGS='-Xmx512m' **Note:** The value 512m in -Xmx512m states that the current configured heap size is 512 MB, you can increase the heap size according to the requirement. 3. Restart the agent.

Workspace troubleshooting

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

Table 8 on page 90 contains problems and solutions related to workspaces.

Table 8. Workspace problems and solutions

Problem	Solution
The process application components are available, but the Availability status shows PROCESS_DATA_NOT_ AVAILABLE.	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: 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 that is on top of the right pane. The Add Counters window is displayed. 4. Look for **Process** in the **Performance object** menu. 5. Complete one of the following actions: • 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.
The name of the attribute does not display in a bar chart or graph view.	When a chart or graph view that includes the attribute is scaled to a small size, a blank space is displayed instead of a truncated name. To see the name of the attribute, expand the view of the chart until sufficient space is available to display all characters of the attribute name.
At the bottom of each view, you see the following Historical workspace KFWITM220E error: `Request failed during execution`.	Ensure that you configure all groups that supply data to the view. In the Historical Configuration view, ensure that data collection is started for all groups that supply data to the view.
You start collection of historical data but the data cannot be seen.	Use the following managing options for historical data collection: • Basic historical data collection populates the Warehouse with raw data. This type of data collection is turned off by default. See the *IBM Tivoli Monitoring Administrator's Guide* for information about managing this feature including how to set the interval at which data is collected. By setting a more frequent interval for data collection, you reduce the load on the system incurred every time data is uploaded. • 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. See the *IBM Tivoli Monitoring Administrator's Guide* to learn how to modify the default collection settings.

Table 8. Workspace problems and solutions (continued)

Problem	Solution
Historical data collection is unavailable because of incorrect queries in the Tivoli Enterprise Portal.	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. Even if data collection has been 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). To ensure support of historical data collection, do not use the Sort By, Group By, or First/Last functions in your queries. See the *IBM Tivoli Monitoring Administrator's Guide* or the Tivoli Enterprise Portal online help for information about the Historical Data Collection function.
When you use a long process name in the situation, the process name is truncated.	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.
Regular (non-historical) monitoring data fails to be displayed.	Check the formation of the queries you use to gather data. For example, look for invalid SQL statements.
No row of data for 64-bit applications is displayed in the workspaces when the monitoring agent is running on a 64-bit operating system.	The Tivoli Enterprise Portal shows data only for 32-bit applications. No solution is available for this problem at this time.
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.	Ensure that application support has been added on the monitoring server, portal server, and portal client. For more information about installing application support, see "Installing and enabling application support" in the *IBM Tivoli Monitoring Installation and Setup Guide*.

Situation troubleshooting

Problems can occur with situations and situation configuration.

Table 9 contains problems and solutions for situations.

Table 9. Situation problems and solutions

Problem	Solution
Monitoring activity requires too much disk space.	Check the RAS trace logging settings that are described in "Setting RAS trace parameters by using the GUI" on page 78. 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 59 describes the performance impact of specific attribute groups. If possible, decrease your use of the attribute groups that require greater system resources.

Table 9. Situation problems and solutions (continued)

Problem	Solution
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: `LT #'Linux_VM_Stats.Total_Memory' / 10`	This formula is incorrect because situation predicates support only logical operators. Your formulas cannot have mathematical operators. **Note:** The Situation Editor provides alternatives to math operators. 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.	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 in the lower right of the window to set the status and appearance of the Situation when it triggers. **Note:** The **State** setting is not related to severity settings in IBM 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.	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 `KpcENV` file for the agent (where *pc* is the two-letter product code): `CDP_NT_EVENT_LOG_CACHE_TIMEOUT=3600` This variable determines how long events from the NT Event Log are kept.
If the Expert Advice for a situation contains a hyperlink to an external website (for example, a Microsoft TechNet website) and you click the hyperlink, the website opens in an external window. However, the external window stops responding.	The external window responds after you close the Preview window and Situation Editor window.
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 *IBM Tivoli Monitoring Installation and Setup Guide*.
The monitoring interval is too long.	Access the Situation Editor view for the situation that you want to modify. Check the **Sampling interval** area in the **Formula** tab. Adjust the time interval as required.
The situation did not activate at startup.	Manually recycle the situation as follows: 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.
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.

Table 9. Situation problems and solutions (continued)

Problem	Solution
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.	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. In the **Formula** tab, analyze predicates as follows: 1. Click the **fx** icon in the upper-right corner of the **Formula** area. The Show formula window is displayed. a. Confirm the following details in the **Formula** area at the top of the window: • The attributes that you intend to monitor are specified in the formula. • The situations that you intend to monitor are specified in the formula. • The logical operators in the formula match your monitoring goal. • The numeric values in the formula match your monitoring goal. b. (Optional) Select the **Show detailed formula** check box in the lower left of the window to see the original names of attributes in the application or operating system that you are monitoring. c. Click **OK** to dismiss the Show formula window. 2. (Optional) In the **Formula** area of the **Formula** tab, temporarily assign numeric values that immediately trigger a monitoring event. The triggering of the event confirms that other predicates in the formula are valid. **Note:** After you complete this test, you must restore the numeric values to valid levels so that you do not generate excessive monitoring data based on your temporary settings. See the *IBM Tivoli Monitoring Troubleshooting Guide* for additional information about situations that do not fire.
Situation events are not displayed in the Events Console view of the workspace.	Associate the situation with a Navigator item. **Note:** The situation does not need to be displayed in the workspace. It is sufficient that the situation is associated with any Navigator item.

Table 9. Situation problems and solutions (continued)

Problem	Solution
You do not have access to a situation.	**Note:** You must have administrator privileges to complete these steps. 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.	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.

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 and follow the instructions.
- Go to the IBM Tivoli Distributed Monitoring and Application Management Wiki. 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.

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 NetApp Storage agent events, the event classes correspond to the agent attribute groups, and the agent-specific slots correspond to the attributes in the attribute group.

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

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

Each of the event classes is a child of KNU_Base and is defined in the `knu.baroc` (version 07.20.03) file. The KNU_Base event class can be used for generic rules processing for any event from the IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage.

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

- aggregate_key: STRING
- aggregate_name: STRING
- aggregate_state: STRING
- aggregate_type: STRING
- avail_size: INTEGER
- block_rate: INTEGER
- block_type: STRING
- cluster_key: STRING
- cp_reads: INTEGER
- cp_read_blocks: INTEGER
- cp_read_latency: INTEGER
- file_rate: INTEGER
- latency: INTEGER
- node: STRING
- node_key: STRING
- percent_overcommit: INTEGER
- raid_size: INTEGER

- raid_status: STRING
- size_available_percent: INTEGER
- size_used_percent: INTEGER
- snapshot_size_available: INTEGER
- snapshot_size_used: INTEGER
- throughput: INTEGER
- timestamp: STRING
- total_committed: INTEGER
- total_size: INTEGER
- total_transfers: INTEGER
- userreadlatency: INTEGER
- userwritelatency: INTEGER
- user_reads: INTEGER
- user_read_blocks: INTEGER
- user_writes: INTEGER
- user_write_blocks: INTEGER
- utilization: INTEGER

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

- average_processor_busy: INTEGER
- cluster_key: STRING
- cluster_node_name: STRING
- cpu_busy: INTEGER
- failed_fan_count: INTEGER
- failed_power_supply_count: INTEGER
- failover_state: STRING
- is_node_healthy: STRING
- max_aggregate_utilization: INTEGER
- memory_size: INTEGER
- net_data_recv: INTEGER
- net_data_sent: INTEGER
- node: STRING
- node_key: STRING
- number_of_processors: INTEGER
- nvram_battery_status: STRING
- other_ops: INTEGER
- read_ops: INTEGER
- read_throughput: INTEGER
- serial_number: STRING
- sys_avg_latency: INTEGER
- sys_read_latency: INTEGER
- sys_write_latency: INTEGER
- timestamp: STRING
- total_throughput: INTEGER

- utilization: INTEGER
- write_throughput: INTEGER

For events that are generated by situations in the Clusters attribute group, events are sent by using the ITM_KNU_CLUSTERS event class. This event class contains the following slots:
- cluster_key: STRING
- cluster_name: STRING
- is_metro_cluster: INTEGER
- knu_status: STRING
- location: STRING
- management_ip: STRING
- max_aggregate_utilization: INTEGER
- max_node_utilization: INTEGER
- metro_cluster_mode: STRING
- node: STRING
- timestamp: STRING
- total_ops: INTEGER
- total_throughput: INTEGER
- version: STRING

For events that are generated by situations in the DataSource attribute group, events are sent by using the ITM_KNU_DATASOURCE event class. This event class contains the following slots:
- agent_connection: STRING
- data_source: STRING
- ip_address: STRING
- node: STRING
- port: STRING
- refresh_interval: STRING
- timestamp: STRING
- version: STRING

For events that are generated by situations in the Disks attribute group, events are sent by using the ITM_KNU_DISKS event class. This event class contains the following slots:
- block_rate: INTEGER
- cluster_key: STRING
- composite_latency: INTEGER
- container_type: STRING
- cp_reads: INTEGER
- cp_read_blocks: INTEGER
- cp_read_latency: INTEGER
- disk_busy: INTEGER
- disk_key: STRING
- disk_name: STRING
- failed_reason: STRING
- file_rate: INTEGER
- guaranteed_read_latency: INTEGER
- guaranteed_write_latency: INTEGER

- is_offline: STRING
- is_virtual: STRING
- node: STRING
- serial_number: STRING
- timestamp: STRING
- total_bytes: INTEGER
- total_latency: INTEGER
- total_ops: INTEGER
- total_throughput: INTEGER
- total_transfers: INTEGER
- used_bytes: INTEGER
- user_reads: INTEGER
- user_read_blocks: INTEGER
- user_read_latency: INTEGER
- user_writes: INTEGER
- user_write_blocks: INTEGER
- user_write_latency: INTEGER

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

- event_key: STRING
- event_severity: STRING
- event_source_key: STRING
- event_status: STRING
- event_time: STRING
- event_type: STRING
- node: STRING
- timestamp: STRING

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

- critical: INTEGER
- entity: STRING
- error: INTEGER
- information: INTEGER
- node: STRING
- normal: INTEGER
- timestamp: STRING
- warning: INTEGER

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

- avg_latency: INTEGER
- is_mapped: STRING
- is_online: STRING
- lun_key: STRING
- lun_path: STRING

- lun_size: INTEGER
- node: STRING
- other_ops: INTEGER
- qtree_key: STRING
- read_data: INTEGER
- read_ops: INTEGER
- serial_number: STRING
- size_used: INTEGER
- storagevm_key: STRING
- timestamp: STRING
- total_ops: INTEGER
- volume_key: STRING
- write_data: INTEGER
- write_ops: INTEGER

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

- error_code: INTEGER
- node: STRING
- object_name: STRING
- object_status: INTEGER
- object_type: INTEGER
- query_name: STRING
- timestamp: STRING

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

- disk_limit: INTEGER
- disk_used: INTEGER
- file_limit: INTEGER
- file_used: INTEGER
- knu_status: STRING
- node: STRING
- qtree_key: STRING
- qtree_name: STRING
- security_style: STRING
- soft_file_limit: INTEGER
- storagevm_key: STRING
- timestamp: STRING
- volume_key: STRING

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

- aggregate_key: STRING
- average_latency: INTEGER
- cifs_other_latency: INTEGER

- cifs_other_ops: INTEGER
- cifs_read_latency: INTEGER
- cifs_read_ops: INTEGER
- cifs_write_latency: INTEGER
- cifs_write_ops: INTEGER
- inode_files_total: INTEGER
- inode_files_used: INTEGER
- iscsi_read_ops: INTEGER
- nfs_other_latency: INTEGER
- nfs_other_ops: INTEGER
- nfs_read_latency: INTEGER
- nfs_read_ops: INTEGER
- nfs_write_latency: INTEGER
- nfs_write_ops: INTEGER
- node: STRING
- other_latency: INTEGER
- other_ops: INTEGER
- overwrite_reserve_actual_used: INTEGER
- percentage_snapshot_reserve: INTEGER
- percentage_snapshot_reserve_used: INTEGER
- quota_over_committed: INTEGER
- quota_status: STRING
- read_latency: INTEGER
- read_ops: INTEGER
- read_throughput: INTEGER
- sis_status: STRING
- size_avail_percent: INTEGER
- size_total: INTEGER
- size_used: INTEGER
- size_used_by_sanpshots: INTEGER
- size_used_percent: INTEGER
- storagevm_key: STRING
- timestamp: STRING
- total_ops: INTEGER
- total_throughput: INTEGER
- volume_key: STRING
- volume_name: STRING
- volume_size: INTEGER
- volume_state: STRING
- volume_style: STRING
- write_latency: INTEGER
- write_ops: INTEGER
- write_throughput: INTEGER

Appendix B. Documentation library

A variety of publications are relevant to the use of the IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage.

The *Using the publications*, SC23-8816, contains information about accessing and using publications. You can find the Documentation Guide in the following information centers:

- IBM Tivoli Monitoring and OMEGAMON® XE
- IBM Tivoli Composite Application Manager

To open the Documentation Guide in the information center, select **Using the publications** in the **Contents** pane.

To find a list of new and changed publications, click **What's new in the information center** on the Welcome page of the IBM Tivoli Monitoring and OMEGAMON XE Information Center.

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

IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage library

The documentation for this agent and other product components is located in the IBM Tivoli Monitoring for Virtual Environments knowledge center.

One document is specific to the NetApp Storage agent: IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage User's Guide. This publication provides agent-specific information for configuring, using, and troubleshooting the NetApp Storage agent.

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

Use the information in this guide with the *Tivoli Enterprise Portal User's Guide* to monitor NetApp 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 Readme First*
- *Exploring IBM Tivoli Monitoring*
- *IBM Tivoli Monitoring Administrator's Guide*
- *IBM Tivoli Monitoring Agent Builder User's Guide*
- *IBM Tivoli Monitoring Command Reference*
- *Configuring IBM Tivoli Enterprise Monitoring Server on z/OS*
- *IBM Tivoli Monitoring Installation and Setup Guide*
- *IBM Tivoli Monitoring: Messages*
- *Using the publications*
- *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: Upgrading from Tivoli Distributed Monitoring*
- *IBM Tivoli Monitoring: Upgrading from V5.1.2*
- *IBM Tivoli Monitoring: i5/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*

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:

- IBM Tivoli Monitoring
- IBM Tivoli Netcool/OMNIbus
- IBM Tivoli Application Dependency Discovery Manager (TADDM)
- IBM 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:

- IBM Integrated Service Management Library is an online catalog that contains integration documentation as well as other downloadable product extensions.
- IBM Redbook publications (Redbooks® publications, Redpapers, and Redbooks technotes) provide information about products from platform and solution perspectives.
- Technotes, which are found through the IBM Software Support website, provide the latest information about known product limitations and workarounds.
- Tivoli wikis

  Tivoli Wiki Central is the home for interactive wikis that offer best practices and scenarios for using Tivoli products. The wikis contain white papers contributed by IBM employees, and content created by customers and business partners.

  Two of these wikis are of particular relevance to IBM Tivoli Monitoring:

  – Tivoli Distributed Monitoring and Application Management Wiki provides information about IBM Tivoli Monitoring and related distributed products, including IBM Tivoli Composite Application Management products.
  – Tivoli System z® Monitoring and Application Management Wiki provides information about the OMEGAMON XE products, Tivoli NetView® for z/OS®, Tivoli Monitoring Agent for z/TPF, and other System z monitoring and application management products.

Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully.

The major accessibility features in this product enable users in the following ways:

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

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

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

Navigating the interface using the keyboard

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

Magnifying what is displayed on the screen

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

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-IBMproducts. 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.

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

B

C

D

E

F

G

H

I

K

L

M

N

O

P

Q

R

S

T

U

V

W

Printed in USA

SC14-7489-02