IBM Tivoli Monitoring for Virtual Environments Agent for VMware VI 7.3 Fix Pack 3

Reference

Note

Before using this information and the product it supports, read the information in "Notices" on page 263.

This edition applies to version 7.3.0.3 of IBM® Tivoli® Monitoring for Virtual Environments Agent for VMware VI (product number 5724-L92) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Tables

Chapter 1. Workspaces

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

About workspaces

Use the Navigator 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, 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.

The IBM Tivoli Monitoring for Virtual Environments Agent for VMware VI provides various default workspaces. These workspaces are displayed in the Navigator under the following nodes and subnodes for this monitoring agent:

VMware VI
Corresponds to a VMware VI instance and contains agent instance-level workspaces.

VMware VI
Each subnode is an ESX server.

When multiple instances of the monitoring agent are defined on a system, the top-level node becomes VMware VI. The VMware VI workspace is undefined at this node. A node for each instance is created called *Instance*::VM. A workspace that is called *Instance*::VM is associated with the instance node. This workspace is comparable to the VMware VI workspace.

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 15.

If you are using remote management to navigate to your systems in the Tivoli Enterprise Portal, navigate from the host name of the computer where you installed the agent.

Predefined workspaces

The VMware VI agent provides predefined workspaces, which are organized by navigator item.

Agent-level navigator items

- VMware VI navigator item
  - VMware VI workspace
  - IBM Systems Director workspace
  - Virtual Enterprise workspace
- Clusters navigator item
  - Cluster Detail workspace
  - Cluster Performance workspace
  - Cluster Summary workspace
  - Clusters workspace
  - Distributed Resource Scheduler workspace
  - Virtual App workspace
- Datastores navigator item
  - Datastore and Volumes workspace
  - Datastore Detail - NAS workspace
  - Datastore Detail - VMFS workspace
  - Datastores workspace
  - Topology - Datastore workspace
  - Virtual Machines Topology workspace
  - VM Datastore Utilization workspace
  - VM Orphaned Disk workspace
- Events navigator item
  - Events workspace
  - Triggered Alarms workspace
- Monitored Servers navigator item
  - Monitored Servers workspace
  - Topology - Monitored Servers workspace
  - Virtual Machines - Monitored Servers workspace
- Networks navigator item
  - Distributed Network Detail workspace
  - Distributed Virtual Switch Detail workspace
  - Network Detail workspace
  - Network NIC Detail workspace
  - Networks workspace

VMware VI (ESX) subnode

- VMware VI navigator item
  - VMware VI workspace
  - All Orphaned Virtual Machines workspace
  - All Virtual Machines workspace

- CPU navigator item
  - CPU workspace
- Disk navigator item
  - Disk workspace
  - Server Disk Detail workspace
  - Server Disk Performance workspace
- ESX Server navigator item
  - Agent Health workspace
  - ESX Server workspace
  - Server DataStore workspace
  - Server Health workspace
- Memory navigator item
  - Memory workspace
- Network navigator item
  - Network workspace
- Resource Pools navigator item
  - Resource Pools workspace
- Virtual Machines navigator item
  - Virtual Machines workspace

Workspace descriptions

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

Workspaces are listed under navigator items. When the agent has subnodes, the navigator items are listed under the subnode.

VMware VI navigator item

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

VMware VI workspace
This workspace provides a snapshot of the health of clusters and data stores. Key indicators show the status to aid in problem identification.

This workspace contains the following views:

Clusters
This view contains key status and performance indicators for clusters. A link is provided to navigate to a workspace with more detailed information about this cluster.

Datastores
This view contains key status and performance indicators for data stores. A link is provided to navigate to a workspace with more detailed information about the data store.

Networks
This view displays all of the configured networks by data center and provides a summary of the health of the network. A link is provided to view the triggered alarms by network.

IBM Systems Director workspace
This workspace provides the IBM Systems Director Web UI to the Director Server this agent is configured to use. It is only available as a workspace link target.

This workspace contains the following view:

IBM Systems Director
This view contains the IBM Systems Director Server Web interface.

Virtual Enterprise workspace
This workspace provides high-level views of the ESX servers that this agent is monitoring.

This workspace contains the following views:

Virtual Center Events
This view contains events that were generated by a monitoring data source. The events are typically specific to the data source.

Monitored Servers
This view shows the ESX servers that this agent is actively monitoring. Each entry in this view contains a link that provides navigation to the ESX Servers and IBM Systems Director workspaces. The IBM Systems Director workspaces require that an IBM Systems Director Server has been configured for the agent and shows the ESX Server in the IBM Systems Director Web UI.

Clusters navigator item

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

Cluster Detail workspace
This workspace contains views that are specific to one cluster. The metrics in the workspace are the detailed metrics of the cluster. The metrics include metrics that are configuration settings and metrics that represent a snapshot of some key performance metrics. Links to other workspaces provided by this agent are included in this workspace.

This workspace contains the following views:

Cluster_name - Datacenter_Name
This view contains a summary of memory and CPU usage for the selected cluster and an overall picture of the health of the cluster.

CPU Utilization - Cluster_Name - Datacenter_Name
This view contains a graph of the number of hosts operating within CPU usage ranges. This view allows a capacity planner or administrator to see how well the CPU resources of the cluster are being used across the entire cluster.

Memory Utilization - Cluster_Name - Datacenter_Name
This view contains a graph of the number of hosts operating within Memory usage ranges. This view allows a capacity planner or administrator to see how well the memory resources of the cluster are being used across the entire cluster.

Navigator
This view contains a navigation aid to quickly jump to views about the other known clusters.

Cluster Performance workspace
This workspace contains views that are specific to one cluster.

This workspace contains the following views:

vMotions vs VMs Powered On for Cluster - Cluster_Name - Datacenter_Name
This view uses a line graph over time to show the number of virtual machines in the given cluster that are powered on. The view also shows the number of virtual machines that have migrated. Historical data collection must be enabled for this view to contain data. See the IBM Tivoli Monitoring Administrator's Guide for details about how to create historical collection. A collection must be created for the Clusters attribute group.

CPU vs Memory Utilization for Cluster - Cluster_Name - Datacenter_Name
This view uses a line graph over time to show the usage of cluster resources CPU and memory. Historical data collection must be enabled for this view to contain data. See the IBM Tivoli Monitoring Administrator's Guide for details about how to create a historical collection. A collection must be created for the Clusters attribute group.

Navigator
This view contains a navigation aid to quickly jump to views about the other known clusters.

Cluster Summary workspace
This workspace contains views that are specific to one cluster. The views in this workspace provide a quick guide to all of the ESX servers, resource pools, and virtual machines that are contained within this cluster. Links are provided to quickly jump to a specific view.

This workspace contains the following views:

ESX Servers - Cluster_Name - Datacenter_Name
This view contains a list of the ESX servers that are members of this cluster. Basic performance data is shown for each server. By selecting the link within this view, you can quickly navigate to the ESX server view depicted in the row of data. The resulting workspace aids in providing additional detailed metrics regarding the ESX server. You can easily navigate back to this Cluster Summary workspace by selecting the appropriate icon.

Resource Pools - Cluster_Name - Datacenter_Name
This view contains a list of the resource pools that are the members of this cluster and shows the basic performance data for each resource pool. This view also provides a link to quickly navigate to the Virtual App workspace, and the link is available only if the Node Type is kvm.Virtual_App.

Datastores - Cluster_Name - Datacenter_Name
This view contains a list of the data stores that are members of this cluster. Basic performance data is shown for each data store.

Virtual Machines - Cluster_Name - Datacenter_Name
This view contains a list of the virtual machines that are members of this cluster. Basic performance data is shown for each virtual machine. By selecting the link within this view, you can quickly navigate to the virtual machine view specific to this virtual machine. The resulting workspace aids in providing additional detailed metrics regarding the virtual machine. You can easily navigate back to this Cluster Summary workspace by selecting the appropriate icon.

Navigator
This view contains a navigation aid to quickly jump to views about the other known clusters.

Clusters workspace
This workspace provides a snapshot of the defined clusters.

This workspace contains the following views:

Clusters
This view contains a list of all of the clusters and a summary of memory and CPU usage for each cluster. Each entry in this view contains a link that provides the ability to navigate to the Cluster Summary, Cluster Details, Distributed Resource Scheduler and IBM Systems Director workspaces. The latter requires that an IBM Systems Director Server has been configured for the agent and shows the Cluster in Topology Common view.

Top 5 by CPU
This view contains a list of the clusters that are ordered by CPU usage.

Top 5 by Memory
This view contains a list of the clusters that are ordered by memory usage.

Bottom 5 by CPU
This view contains a list of the clusters that are ordered by CPU usage.

Bottom 5 by Memory
This view contains a list of the clusters that are ordered by memory usage.

Distributed Resource Scheduler workspace
This workspace contains view that is specific to the Distributed Resource Scheduler (DRS) and Storage Distributed Resource Scheduler (SDRS) faults.

This workspace contains the following view:

DRS Faults
This view shows information about the DRS and SDRS faults that are associated with the clusters.

Virtual App workspace
This workspace contains views that are specific to the virtual application.

This workspace contains the following view:

Virtual App
This view shows information about the virtual machines and virtual applications that are associated with the virtual machines.

Datastores navigator item

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

Datastore and Volumes workspace
This workspace contains views that are specific to one data store. The metrics are supplied by both the IBM Tivoli Monitoring for Virtual Environments Agent for VMware VI and the Agent for NetApp Storage. This workspace provides for both a virtual and a physical view of the data store. Use the Storage Agent tab on the configuration windows to set up the relationship between the agents.

This workspace contains the following views:

Datastore Health
This view shows configuration metrics from the virtualized environment. This view is primarily configuration data.

Volume by Operations
This view shows physical metrics about the I/O operations and is supplied by the Agent for NetApp Storage.

Volume by Latency
This view describes the latency of the data store on a physical volume.

Volume by Transfer Rate
This view describes the I/O transfer rates for this physical volume.

Datastore Detail - NAS workspace
This workspace contains views that are specific to one data store. The metrics in the workspace are the detailed metrics of the data store. The metrics include metrics that are configuration settings and metrics that represent a snapshot of some key performance metrics.

This workspace contains the following views:

Utilization
This view displays graphically the usage percentage of this data store.

Connections
This view shows the dependencies of other virtualization components on this data store.

Percent Used - History
This view shows percentage used of this data store over time. The time period is configurable. This data helps identify trends and spikes that occur at various points in time. Historical collection must be enabled for this view to populate.

Datastore Detail
This view contains the detailed configuration specifications of this data store and additional usage metrics.

Volumes
This view contains data when an additional IBM Tivoli Monitoring agent has been configured and the data store is located on a NetApp or IBM Series N storage device.

Topology
This view is a link to topology workspaces that have this data store as a node. The status of the data store is depicted by the icon.

Datastore Detail - VMFS workspace

This workspace contains views that are specific to one data store. The metrics in the workspace are the detailed metrics of the data store. The metrics include metrics that are configuration settings and metrics that represent a snapshot of some key performance metrics.

This workspace contains the following views:

Utilization

This view displays graphically the usage percentage of this data store.

Connections

This view shows the dependencies of other virtualization components on this data store.

Percent Used - History

This view shows percentage used of this data store over time. The time period is configurable. This data helps identify trends and spikes that occur at various points in time. Historical collection must be enabled for this view to populate.

Datastore Detail

This view contains the detailed configuration specifications of this data store and additional usage metrics.

Topology

This view is a link to topology workspaces that have this data store as a node. The status of the data store is depicted by the icon.

Datastores workspace

This workspace contains a list of all the data stores. This list might be used to identify problems with the data store. More detailed information about a specific data store can be obtained by using the link next to a row describing a data store.

This workspace contains the following views:

Datastore Health

This view shows all the data stores and basic health indicators for each one. In addition, information regarding how many other components are connected to the data store is shown. This information is helpful in providing insight about the impact of performance problems that the data store might be experiencing.

NAS Datastores

This view is specific to all data stores that are not of the VMFS type. This view is typically data stores backed by network-attached devices and defined on NFS or CIFS volumes. The link on each row enables navigation to a more detailed workspace specific to that data store.

VMFS Datastores

This view is specific to all data stores that are of the VMFS type. Data stores of type VMFS can be local to an ESX host or attached through a SAN device. The link on each row enables navigation to a more detailed workspace specific to that data store.

Datastore Clusters

This view shows all the data store clusters. The link in each row enables the navigation to a workspace that is specific to the data store cluster.

Topology - Datastore workspace

This workspace shows the relationship between data stores and ESX servers and clusters.

This workspace contains the following view:

Topology

This view displays graphically the logical connections of the data stores to ESX servers and clusters. The status of each entity is also depicted by each icon.

Virtual Machines Topology workspace

This workspace shows the relationship between virtual machines and other entities in the virtual enterprise such as data stores and clusters.

This workspace contains the following view:

Topology
This view displays graphically the logical connections of the virtual machines to ESX servers, data stores, and clusters. The status of each entity is also depicted by each icon.

VM Datastore Utilization workspace
This workspace contains views that are specific to one data store. These metrics provide insight about which virtual machines are allocated to this data store.

This workspace contains the following views:

VM Datastore Utilization
This view displays metrics that illustrate which virtual machines are allocated on this data store. These metrics show how much space the virtual machine is currently using and how much space the virtual machine is allowed to use as it grows.

VM IO Operations
This view displays the amount of data being read and written by the virtual machines on this datastore. These metrics show how busy the datastore is by virtual machine.

VM Datastore Provisioned Space
This view shows graphically how much of the total provisioned space that is given to the virtual machine is actually being used.

Total IO by VM
This view shows graphically how total data from both read and write operations is being done by each virtual machine.

VM Orphaned Disk workspace
This workspace displays details about the orphaned virtual machine disk.

This workspace contains the following view:

VM Orphaned Disk
This view shows the details about the space that is used by an orphaned virtual machine disk on the data store. In addition, this view shows the date and time when an orphaned virtual machine disk was last modified.

Events navigator item

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

Events workspace
This workspace contains a list of events that have occurred while the monitoring agent is running. The events that are listed are not specific to an ESX server, but they are specific to a configured data source.

This workspace contains the following views:

Virtual Center Events
This view contains a list of the events specific to the data source.

Triggered Alarms
This view contains a list of the alarms that are triggered by VMWare for various monitored entities such as data stores and ESX hosts.

Virtual Center Tasks
This view contains a list of the tasks that are triggered on the vCenter server, and the tasks that are completed or failed for various monitored entities, such as data stores and the ESX hosts.

Virtual Center Active Tasks
This view contains a list of the active tasks that are triggered on the vCenter server for various monitored entities, such as data stores and the ESX hosts.

Triggered Alarms workspace
This workspace contains a view that lists the alarms for a specific entity in order to be able to quickly identify the alarms that influence the alarms that have triggered on that object.

This workspace contains the following view:

Triggered Alarms
This view contains a list of the alarms that are triggered by VMWare for various monitored entities such as data stores and ESX hosts.

Monitored Servers navigator item

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

Monitored Servers workspace
This workspace contains a list of the monitored ESX servers.

This workspace contains the following views:

Monitored Servers
This view contains a list of the monitored ESX servers. Monitored servers are discovered from the agent data source, which can be a VMware Virtual Center or an ESX server. Each entry in this view contains a link that provides navigation to the ESX Servers and IBM Systems Director workspaces. The IBM Systems Director workspaces require that an IBM Systems Director Server has been configured for the agent and shows the ESX Server in the IBM System Director Web UI.

Data Sources
This view provides status information about the data sources that the agent uses to collect monitoring data.

Agent Events
This view provides status information about the agent that is helpful if there is a configuration issue or if there is a problem connecting to a vCenter or ESX server.

Topology - Monitored Servers workspace
This workspace provides insight into the logical connections between the major entities in the virtual enterprise.

This workspace contains the following view:

Topology
This view shows the relationship among virtual machines, ESX servers, clusters, resource pools, data centers, and the vCenter.

Virtual Machines - Monitored Servers workspace
This workspace shows the relationship between virtual machines and other entities in the virtual enterprise such as ESX servers and clusters.

This workspace contains the following view:

Topology
This view displays graphically the logical connections of the virtual machines to ESX servers, and clusters. The status of each entity is also depicted by each icon.

Networks navigator item

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

Distributed Network Detail workspace
This workspace provides detail of a selected network in the infrastructure.

This workspace contains the following views:

Network - Network_Name
This view displays the selected network status and configuration. A link is provided to view the triggered alarms by network.

Networked Virtual Machines - Network_Name
This view shows the networked Virtual Machines usage.

Networked Servers - Network_Name
This view shows the networked servers usage.

Distributed Virtual Switches - Switch_Name
This view shows the Distributed virtual switches usage.

Distributed Virtual Switch Detail workspace
This workspace provides detail of a selected Distributed Virtual switch in the infrastructure.

This workspace contains the following views:

Distributed Virtual Switch - Switch_Name
This view displays the selected Distributed Virtual switch in the virtual infrastructure.

Distributed Virtual Uplinks - Switch_Name
This view displays all the distributed virtual uplinks that are associated with the selected Distributed Virtual switch.

Distributed Virtual Portgroups - Switch_Name
This view displays all the distributed virtual port groups that are associated with the selected Distributed Virtual switch.

DVS Host Member Health - Switch_Name
This view displays health status of all the host associated with the selected Distributed Virtual switch.

Network Detail workspace
This workspace provides detail of a selected network in the infrastructure.

This workspace contains the following views:

Network - Network_Name
This view displays the selected network status and configuration. A link is provided to view the triggered alarms by network.

Networked Virtual Machines - Network_Name
This view shows the networked Virtual Machines usage.

Networked Servers - Network_Name
This view shows the networked servers usage.

Networked Virtual Switches - Network_Name
This view shows the networked virtual switches usage.

Network NIC Detail workspace
This workspace provides detail information about Network NIC in the infrastructure.

This workspace contains the following views:

Networked Virtual Machines - Switch_Name
This view displays the Networked Virtual Machines that are associated with the selected switch.

Distributed Virtual Uplinks - Switch_Name - Host_Name
This view displays all the distributed virtual uplinks that are associated with the selected Distributed Virtual switch.

Distributed Virtual Uplinks - Switch_Name - Portgroup_Name
This view displays all the distributed virtual port groups that are associated with the selected Distributed Virtual switch.

Networks workspace
This workspace displays a summary of all the networks that are configured within the data centers.

This workspace contains the following views:

Networks
This view displays all of the configured networks by data center and provides a summary of the health of the network. A link is provided to view the triggered alarms by network.

Standard Virtual Switches
This view displays all of the virtual standard switches in the virtual infrastructure.

Distributed Virtual Switches
This view displays all of the Distributed Virtual switches in the virtual infrastructure.

VMware VI subnode

The predefined workspace descriptions for the subnode are organized by the navigator item to which the workspaces are relevant.

VMware VI navigator item

VMware VI workspace
This workspace provides views that show performance indicators for a single ESX server or host.

This workspace contains the following views:

Server CPU Utilization
This view shows CPU usage of the server or host by individual CPU.

Server Memory Utilization
This view shows the overall memory usage of the server.

Server Network
This view shows the network performance of the server by network interface.

All Orphaned Virtual Machines workspace
This workspace shows details about all the orphaned virtual machines.

This workspace contains the following view:

All Orphaned Virtual Machines
This view contains a list of the orphaned virtual machines of virtual environment.

All Virtual Machines workspace
This workspace shows the details of all the virtual machines.

This workspace contains the following views:

Virtual Machines
This view shows the details about status of the virtual machines. In addition, this view contains a list of the virtual machines.

Virtual Machines by CPU
This view contains a list of the virtual machines that are categorized by CPU.

Virtual Machines by Memory
This view contains a list of the virtual machines that are categorized by memory.

CPU navigator item

CPU workspace
This workspace provides views of the CPU usage of the ESX server.

This workspace contains the following views:

Virtual Machine CPU
This view shows the CPU usage of the virtual machines on this ESX server that are powered on.

Utilization by Virtual Machine Name - CPU%
This view shows the CPU usage of each virtual machine, by CPU, that is powered on.

Percent Ready by Virtual Machine Name - CPU
This view shows the CPU Percent Ready attribute for each virtual machine that is powered on. Ideally, this value is low.

Server CPU Percent Usage
This view shows the CPU usage of the ESX server.

CPU Percent Use Per VM
This view shows the CPU usage of the Virtual Machine.

Disk navigator item

Disk workspace
This workspace provides views of the disk usage of the ESX server.

This workspace contains the following views:

Server Disk
This view shows the ESX server disk usage.

Virtual Machine Partitions
This view shows the disk partitions within the virtual machines. Partition information is available only if the virtual machine has the VMware Tools package installed and running.

Virtual Machine Disks
This view shows the virtual disks defined for the virtual machine.

Virtual Machine Disks Performance
This view shows information about the performance of disks that are associated with the virtual machines.

Server Disk Detail workspace
This workspace contains views that are specific to one ESX host. The metrics are for disks from the ESX host point of view. This data includes local disks and data stores visible to this host.

This workspace contains the following views:

Server Disk IO
This graphical view displays metrics for each disk. The metrics show the number of read and write operations on the disk. The metrics command and commands aborted show how well the disk is servicing the requests.

Server Disk Total Latencies
This graphical view displays the total latency values of the disk requests categorized into device, kernel and queue.

Server Disk Details
This view shows a summary of the performance metrics for this disk.

Server Disk Average Latencies
This graphical view displays the average latency metrics for the server disk introduced by the device, kernel and queue.

Server Disk Performance workspace
This workspace contains views that are specific to one ESX host. The metrics are for disks from the ESX host point of view. This data includes local disks and data stores that are visible to this host.

This workspace contains the following views:

Virtual Machine Disks
This view shows the virtual hard disk drives that are configured for the virtual machines that are running on this host.

Server Disk
This view shows performance metrics for the disks defined to this host. This data includes local disks and data stores. These metrics give an idea of the demand on the disk and how well the disk is servicing the requests.

Server SAN
This view shows configuration information about SAN-attached disks on the ESX host.

ESX Server navigator item

Agent Health workspace

This workspace contains a list of the VMware data stores.

This workspace contains the following view:

VMWare VI Agent Status

This view contains a list of the status of the agent data collection operations for the attribute groups.

ESX Server workspace

This workspace provides views that describe the overall operating state of an ESX server. Links that enable quick navigation to other workspaces have been provided in this workspace. To fully take advantage of this feature, the Monitoring Agent for Linux must be installed on the ESX server.

This workspace contains the following views:

Server Summary

This partial view of the ESX server shows the server status and basic information.

Server Parameters

This partial view of the ESX server shows the server status and basic resource consumption.

Events

This view contains a list of events that have recently occurred. VMware alarms and events pertaining to this server are displayed here.

Overall CPU Utilization

This view shows the overall CPU usage of this server. The areas in color represent usage ranges that are noteworthy.

Overall Memory Utilization

This view shows the overall memory usage of this server.

Server DataStore workspace

This workspace contains a list of the VMware data stores.

This workspace contains the following views:

Server DataStore

This view contains a list of the data stores that this server is configured to use.

Server DataStore Usage

This view contains a list of the usage information for each data store.

Server HBAs

This view contains the Host Bus Adaptors for the EXS server.

Server Health workspace

This workspace contains a list of all hardware sensors.

This workspace contains the following view:

Sensors

This view contains a list of all hardware sensors, their status, and their value.

Memory navigator item

Memory workspace

This workspace provides views of the memory usage of the ESX server.

This workspace contains the following views:

Server Memory

This view shows the memory usage of the ESX server.

Virtual Machine Memory
This view shows the memory usage and configuration settings for the virtual machines configured on this ESX server.

Guest Memory Utilization
This view depicts the amount of memory used by the virtual machine guest operating system.

Network navigator item

Network workspace
This workspace provides views of the network usage of the ESX server.

This workspace contains the following views:

Server Network
This view shows the network usage of the ESX server.

Virtual Machine Network
This view shows the network usage of the virtual machines on this ESX server.

Server Virtual Switches
This view shows the virtual switches on this ESX server.

Resource Pools navigator item

Resource Pools workspace
This workspace provides views of the resource pools that are known to the ESX server.

This workspace contains the following views:

Resource Pools
This view shows the general metrics for the resource pools.

Resource Pool CPU
This view shows the CPU metrics for the resource pools.

Resource Pool Memory
This view shows the memory metrics for the resource pools.

Virtual Machines navigator item

Virtual Machines workspace
This workspace provides views of the virtual machines that are defined on this ESX server.

This workspace contains the following views:

Virtual Machines Configuration
This view shows the properties given to virtual machines at the time they were created.

Virtual Machines Status
This view shows some of the runtime metrics for the virtual machines. If the VMware Tools package is not running on the virtual machine or the virtual machine is powered off, then several metrics are not known.

Virtual Machines Snapshots
This view shows information about the snapshots for the virtual machines.

Chapter 2. Attributes

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

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 VMware VI 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” on page 15 and “Attributes in each attribute group” on page 20.

Attribute groups for the monitoring agent

The VMware VI agent contains the following attribute groups. For agents that use IBM Tivoli Monitoring infrastructure, attributes are in attribute groups. For agents that use the lightweight infrastructure, attributes are in data sets.

The table name depends on the maximum table name limits of the target database being used for the historical data collection. If the maximum name is 30 characters, any warehouse or historical table name longer than 30 characters is shortened to 30 characters.

Note : Agents that use the Tivoli Monitoring infrastructure refer to the historical table name as the warehouse table name.

- Attribute group name: Active Tasks
  - Table name: KVMATASKS
  - Warehouse or historical table name: KVM_ACTIVE_TASKS or KVMATASKS
- Attribute group name: Agent Events
  - Table name: KVMAEVENTS
  - Warehouse or historical table name: KVM_AGENT_EVENTS or KVMAEVENTS
- Attribute group name: Cluster DRS Faults

  - Table name: KVMCLTDRSF
  - Warehouse or historical table name: KVM_CLUSTER_DRS_FAULTS or KVMCLTDRSF
- Attribute group name: Clustered Datastores
  - Table name: KVMCLTRDST
  - Warehouse or historical table name: KVM_CLUSTERED_DATASTORES or KVMCLTRDST
- Attribute group name: Clustered Resource Pools
  - Table name: KVMCLTRRPS
  - Warehouse or historical table name: KVM_CLUSTERED_RESOURCE_POOLS or KVMCLTRRPS
- Attribute group name: Clustered Servers
  - Table name: KVMCLTRSRV
  - Warehouse or historical table name: KVM_CLUSTERED_SERVERS or KVMCLTRSRV
- Attribute group name: Clustered Virtual Apps
  - Table name: KVMCLTVAPS
  - Warehouse or historical table name: KVM_CLUSTERED_VIRTUAL_APPS or KVMCLTVAPS
- Attribute group name: Clustered Virtual Machines
  - Table name: KVMCLTRVMS
  - Warehouse or historical table name: KVM_CLUSTERED_VIRTUAL_MACHINES or KVMCLTRVMS
- Attribute group name: Clusters
  - Table name: KVMCLUSTRT
  - Warehouse or historical table name: KVM_CLUSTERS or KVMCLUSTRT
- Attribute group name: Datacenters
  - Table name: KVMDCTRS
  - Warehouse or historical table name: KVM_DATACENTERS or KVMDCTRS
- Attribute group name: Datastore Cluster
  - Table name: KVMDRCLUST
  - Warehouse or historical table name: KVM_DATASTORE_CLUSTER or KVMDRCLUST
- Attribute group name: Datastore Host Disks
  - Table name: KVMDSHSD
  - Warehouse or historical table name: KVM_DATASTORE_HOST_DISKS or KVMDSHSD
- Attribute group name: Datastore Topology
  - Table name: KVMSTOPO
  - Warehouse or historical table name: KVM_DATASTORE_TOPOLOGY or KVMSTOPO
- Attribute group name: Datastores
  - Table name: KVMDSTORES
  - Warehouse or historical table name: KVM_DATASTORES or KVMDSTORES
- Attribute group name: Director
  - Table name: KVMDIRE
  - Warehouse or historical table name: KVM_DIRECTOR or KVMDIRE
- Attribute group name: Distributed Virtual Portgroups
  - Table name: KVMDVPGRPS
  - Warehouse or historical table name: KVM_DISTRIBUTED_VIRTUAL_PORTGROUPS or KVMDVPGRPS
- Attribute group name: Distributed Virtual Switch Health

  - Table name: KVMDVSHLTH
  - Warehouse or historical table name: KVM_DISTRIBUTED_VIRTUAL_SWITCH_HEALTH or KVMDVSHLTH
- Attribute group name: Distributed Virtual Switches
  - Table name: KVMDVSWTCH
  - Warehouse or historical table name: KVM_DISTRIBUTED_VIRTUAL_SWITCHES or KVMDVSWTCH
- Attribute group name: Distributed Virtual Uplinks
  - Table name: KVMDVUPLNK
  - Warehouse or historical table name: KVM_DISTRIBUTED_VIRTUAL_UPLINKS or KVMDVUPLNK
- Attribute group name: ESX Performance Object Status
  - Table name: KVMESXPOS
  - Warehouse or historical table name: KVM_ESX_PERFORMANCE_OBJECT_STATUS or KVMESXPOS
- Attribute group name: Events
  - Table name: KVMIRAEVNT
  - Warehouse or historical table name: KVM_EVENTS or KVMIRAEVNT
- Attribute group name: Monitored Servers
  - Table name: KVMDAG
  - Warehouse or historical table name: KVM_MONITORED_SERVERS or KVMDAG
- Attribute group name: Networked Servers
  - Table name: KVMNETSERV
  - Warehouse or historical table name: KVM_NETWORKED_SERVERS or KVMNETSERV
- Attribute group name: Networked Virtual Machines
  - Table name: KVMNETVM
  - Warehouse or historical table name: KVM_NETWORKED_VIRTUAL_MACHINES or KVMNETVM
- Attribute group name: Networked Virtual Switches
  - Table name: KVMNVSWITC
  - Warehouse or historical table name: KVM_NETWORKED_VIRTUAL_SWITCHES or KVMNVSWITC
- Attribute group name: Networks
  - Table name: KVMDCNETS
  - Warehouse or historical table name: KVM_NETWORKS or KVMDCNETS
- Attribute group name: Performance Object Status
  - Table name: KVMPOBJST
  - Warehouse or historical table name: KVM_PERFORMANCE_OBJECT_STATUS or KVMPOBJST
- Attribute group name: Resource Pool CPU
  - Table name: KVMRSPOOLC
  - Warehouse or historical table name: KVM_RESOURCE_POOL_CPU or KVMRSPOOLC
- Attribute group name: Resource Pool General
  - Table name: KVMRSPOOLG
  - Warehouse or historical table name: KVM_RESOURCE_POOL_GENERAL or KVMRSPOOLG
- Attribute group name: Resource Pool Memory
  - Table name: KVMRSPOOLM
  - Warehouse or historical table name: KVM_RESOURCE_POOL_MEMORY or KVMRSPOOLM

- Attribute group name: Server
  - Table name: KVMSERVERG
  - Warehouse or historical table name: KVM_SERVER or KVMSERVERG
- Attribute group name: Server CPU
  - Table name: KVMSERVERC
  - Warehouse or historical table name: KVM_SERVER_CPU or KVMSERVERC
- Attribute group name: Server DataStore
  - Table name: KVMSERVRDS
  - Warehouse or historical table name: KVM_SERVER_DATASTORE or KVMSERVRDS
- Attribute group name: Server Disk
  - Table name: KVMSERVERD
  - Warehouse or historical table name: KVM_SERVER_DISK or KVMSERVERD
- Attribute group name: Server HBA
  - Table name: KVMSRVHBAS
  - Warehouse or historical table name: KVM_SERVER_HBA or KVMSRVHBAS
- Attribute group name: Server Health
  - Table name: KVMSVRHLTH
  - Warehouse or historical table name: KVM_SERVER_HEALTH or KVMSVRHLTH
- Attribute group name: Server Memory
  - Table name: KVMSERVERM
  - Warehouse or historical table name: KVM_SERVER_MEMORY or KVMSERVERM
- Attribute group name: Server Network
  - Table name: KVMSERVERN
  - Warehouse or historical table name: KVM_SERVER_NETWORK or KVMSERVERN
- Attribute group name: Server SAN
  - Table name: KVMSRVRSAN
  - Warehouse or historical table name: KVM_SERVER_SAN or KVMSRVRSAN
- Attribute group name: Server Virtual Switches
  - Table name: KVMSRVVSWI
  - Warehouse or historical table name: KVM_SERVER_VIRTUAL_SWITCHES or KVMSRVVSWI
- Attribute group name: Server VM Datastore Utilization
  - Table name: KVMSVMDSUT
  - Warehouse or historical table name: KVM_SERVER_VM_DATASTORE_UTILIZATION or KVMSVMDSUT
- Attribute group name: SubNode Events
  - Table name: KVMSERVERE
  - Warehouse or historical table name: KVM_SUBNODE_EVENTS or KVMSERVERE
- Attribute group name: Tasks
  - Table name: KVMTASKS
  - Warehouse or historical table name: KVM_TASKS
- Attribute group name: Thread Pool Status
  - Table name: KVMTHPLST
  - Warehouse or historical table name: KVM_THREAD_POOL_STATUS or KVMTHPLST

- Attribute group name: Topological Events
  - Table name: KVMTOPEVNT
  - Warehouse or historical table name: KVM_TOPOLOGICAL_EVENTS or KVMTOPEVNT
- Attribute group name: Topology
  - Table name: KVMTOPO
  - Warehouse or historical table name: KVM_TOPOLOGY or KVMTOPO
- Attribute group name: Triggered Alarms
  - Table name: KVMALARMS
  - Warehouse or historical table name: KVM_TRIGGERED_ALARMS or KVMALARMS
- Attribute group name: vCenters
  - Table name: KVMVCENTER
  - Warehouse or historical table name: KVM_VCENTERS or KVMVCENTER
- Attribute group name: Virtual Machines
  - Table name: KVMVM_GEN
  - Warehouse or historical table name: KVM_VIRTUAL_MACHINES or KVMVM_GEN
- Attribute group name: Virtual Switches
  - Table name: KVMVSWITCH
  - Warehouse or historical table name: KVM_VIRTUAL_SWITCHES or KVMVSWITCH
- Attribute group name: VM CPU
  - Table name: KVMVM_CPU
  - Warehouse or historical table name: KVM_VM_CPU
- Attribute group name: VM Datastore Utilization
  - Table name: KVMVMDSUTL
  - Warehouse or historical table name: KVM_VM_DATASTORE_UTILIZATION or KVMVMDSUTL
- Attribute group name: VM Disk
  - Table name: KVMVM_DISK
  - Warehouse or historical table name: KVM_VM_DISK
- Attribute group name: VM Disk Performance
  - Table name: KVMVMDKPRF
  - Warehouse or historical table name: KVM_VM_DISK_PERFORMANCE or KVMVMDKPRF
- Attribute group name: VM Memory
  - Table name: KVMVM_MEM
  - Warehouse or historical table name: KVM_VM_MEMORY or KVMVM_MEM
- Attribute group name: VM Network
  - Table name: KVMVM_NET
  - Warehouse or historical table name: KVM_VM_NETWORK or KVMVM_NET
- Attribute group name: VM Orphaned Disk
  - Table name: KVMVMORPDI
  - Warehouse or historical table name: KVM_VM_ORPHANED_DISK or KVMVMORPDI
- Attribute group name: VM Partition
  - Table name: KVMVM_PART
  - Warehouse or historical table name: KVM_VM_PARTITION or KVMVM_PART

- Attribute group name: VM Snapshot
  - Table name: KVMVMSNAP
  - Warehouse or historical table name: KVM_VM_SNAPSHOT or KVMVMSNAP
- Attribute group name: VM SnapshotFileLayout
  - Table name: KVMVMSNPFL
  - Warehouse or historical table name: KVM_VM_SNAPSHOTFILELAYOUT or KVMVMSNPFL
- Attribute group name: VM Snapshots
  - Table name: KVMVMSNAPS
  - Warehouse or historical table name: KVM_VM_SNAPSHOTS or KVMVMSNAPS

Attributes in each attribute group

Attributes in each VMware VI agent attribute group collect data that the agent uses for monitoring.

The description of each attribute group contains the following details:

- Whether the attribute group is a historical type that you can roll off to a data warehouse.
- Information such as whether the attribute is a key attribute, type, source, verification method, warehouse name (as applicable), and other names.

  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.

Active Tasks attribute group

This attribute group provides information about the active tasks that are running on the vCenter server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Cancelable attribute
Indicates whether cancellation of the task is supported. The type is integer with enumerated values. The following values are defined: Yes (1), No (0), Unavailable (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CANCELABLE (warehouse name), Cancelable (caption), Cancelable (attribute name), and CANCELABLE (column name).

Initiated By attribute
The type of the entity that created the task.The valid values are user name, another schedule task name, alarm name, and system. The type is 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.

The following names are defined for this attribute: INITIATED_BY or IB (warehouse name), Initiated By (caption), Initiated_By (attribute name), and IB (column name).

Name attribute
The name of the task. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NAME (warehouse name), Name (caption), Name (attribute name), and NAME (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Queue Time attribute
The date and time when the task was created. The type is 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.

The following names are defined for this attribute: QUEUE_TIME (warehouse name), Queue Time (caption), Queue_Time (attribute name), and QUEUE_TIME (column name).

Source Hostname attribute
The host name of the data source that created the task. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SOURCE_HOSTNAME or SH (warehouse name), Source Hostname (caption), Source_Hostname (attribute name), and SH (column name).

Start Time attribute
The date and time when the task started running. The type is 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.

The following names are defined for this attribute: START_TIME (warehouse name), Start Time (caption), Start_Time (attribute name), and START_TIME (column name).

Status attribute
The current status of the task.The valid values are queued and running. The type is 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.

The following names are defined for this attribute: STATUS (warehouse name), Status (caption), Status (attribute name), and STATUS (column name).

Target Entity attribute
The name of the target managed entity for the task. The type is 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.

The following names are defined for this attribute: TARGET_ENTITY or TE (warehouse name), Target Entity (caption), Target_Entity (attribute name), and TE (column name).

Target Entity Type attribute
The type of the target managed entity for the task. The type is 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.

The following names are defined for this attribute: TARGET_ENTITY_TYPE or TET (warehouse name), Target Entity Type (caption), Target_Entity_Type (attribute name), and TET (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Agent Events attribute group

This attribute group receives messages from the agent about agent status. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Managed System attribute
The managed system that is associated with this event. The type is string.

The following names are defined for this attribute: MANAGED_SYSTEM or MS (warehouse name), Managed System (caption), Managed_System (attribute name), and MS (column name).

Message attribute
The message of this event. The type is integer with enumerated values. The following values are defined: Connection Timeout. (24), Server performance API unavailable. (23), Connection failed: Incorrect WSDL Version (22), Insufficient Permissions (21), Insufficient Permissions: Missing Datastore.Browse (20), Insufficient Permissions: Missing System.Read (19), Insufficient Permissions: Missing System.View (18), Insufficient Permissions: Missing VirtualMachine.Interact.PowerOff (17), Insufficient Permissions: Missing VirtualMachine.Interact.PowerOn (16), Initial Property Collection Complete (15), Connection failed: http redirected (14), Connection failed: unsupported server version (13), Agent Stopped (12), Agent Started (11), Connection failed: unknown failure (10), VM Power On Task Succeeded (9), VM Power On Task Failed (8), VM Power Off Task Succeeded (7), VM Power Off Task Failed (6), Connection reset (5), Connection succeeded (4), Connection failed: username or password invalid (3), Connection failed: ssl negotiation failed (2), Connection failed: connection refused (1), Connection failed: address not found (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MESSAGE (warehouse name), Message (caption), Message (attribute name), and MESSAGE (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Severity attribute
The level of severity for this agent event. The type is integer with enumerated values. The following values are defined: Severe (2), Warning (1), Info (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SEVERITY (warehouse name), Severity (caption), Severity (attribute name), and SEVERITY (column name).

Source attribute
The source of this agent event. The type is string.

The following names are defined for this attribute: SOURCE (warehouse name), Source (caption), Source (attribute name), and SOURCE (column name).

Subsystem attribute
The subsystem of the agent that generated this event. The type is integer with enumerated values. The following values are defined: Permission (2), General (2), Task (1), Connection (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SUBSYSTEM (warehouse name), Subsystem (caption), Subsystem (attribute name), and SUBSYSTEM (column name).

Timestamp attribute
The time the event was generated. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Cluster DRS Faults attribute group

This attribute group provides information about the Distributed Resource Scheduler (DRS) and Storage Distributed Resource Scheduler (SDRS) faults that are generated in the cluster. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Cluster attribute
The name of the cluster where the fault is generated. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: CLUSTER (warehouse name), Cluster (caption), Cluster (attribute name), and CLUSTER (column name).

DataCenter attribute
The name of the data center that contains the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), DataCenter (caption), DataCenter (attribute name), and DATACENTER (column name).

DRS Type attribute
The type of DRS.The valid values are DRS and SDRS. The type is 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.

The following names are defined for this attribute: DRS_TYPE (warehouse name), DRS Type (caption), DRS_Type (attribute name), and DRS_TYPE (column name).

Fault Message attribute
The message that is displayed corresponding to the fault. The type is 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.

The following names are defined for this attribute: FAULT_MESSAGE or FM (warehouse name), Fault Message (caption), Fault_Message (attribute name), and FM (column name).

Fault Name attribute
The name of the fault that is generated in the cluster. The type is 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.

The following names are defined for this attribute: FAULT_NAME (warehouse name), Fault Name (caption), Fault_Name (attribute name), and FAULT_NAME (column name).

FT Virtual Machine attribute
The name of the fault tolerance virtual machine. If this attribute value is Unavailable, the fault is not associated with a particular virtual machine. The type is 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.

The following names are defined for this attribute: FT_VIRTUAL_MACHINE or FVM (warehouse name), FT Virtual Machine (caption), FT_Virtual_Machine (attribute name), and FVM (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Reason attribute
The code that explains why DRS attempted to set recommendations for entities (such as Rule enforcement, Power management, and so on) when faults were generated. The type is 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.

The following names are defined for this attribute: REASON (warehouse name), Reason (caption), Reason (attribute name), and REASON (column name).

Source attribute
The host name of the data source. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SOURCE (warehouse name), Source (caption), Source (attribute name), and SOURCE (column name).

Source Hostname attribute
The name of the host system of a virtual machine.If this attribute value is Unavailable, the fault is not associated with a particular virtual machine. The type is 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.

The following names are defined for this attribute: SOURCE_HOSTNAME or SH (warehouse name), Source Hostname (caption), Source_Hostname (attribute name), and SH (column name).

Target Hostname attribute
The name of the target host system that is selected for the migration of virtual machine. The type is 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.

The following names are defined for this attribute: TARGET_HOSTNAME or TH (warehouse name), Target Hostname (caption), Target_Hostname (attribute name), and TH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Virtual Machine attribute
The name of the virtual machine that the DRS was trying to move when the fault was generated.If this attribute value is Unavailable, the fault is not associated with a particular virtual machine. The type is 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.

The following names are defined for this attribute: VIRTUAL_MACHINE or VM (warehouse name), Virtual Machine (caption), Virtual_Machine (attribute name), and VM (column name).

Clustered Datastores attribute group

This attribute group describes the data stores that are used by a cluster. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Accessible attribute

Whether the data store is accessible or not. The type is integer with enumerated values. The following values are defined: Yes (1), No (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ACCESSIBLE (warehouse name), Accessible (caption), Accessible (attribute name), and ACCESSIBLE (column name).

Capacity attribute

The storage capacity in MB of the data store. The type is 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.

The following names are defined for this attribute: CAPACITY (warehouse name), Capacity (caption), Capacity (attribute name), and CAPACITY (column name).

Cluster attribute

The name of the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: CLUSTER (warehouse name), Cluster (caption), Cluster (attribute name), and CLUSTER (column name).

Cluster MORef attribute

The internal managed object reference name of this cluster. The type is 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.

The following names are defined for this attribute: CLUSTER_MOREF or CM (warehouse name), Cluster MORef (caption), Cluster_MORef (attribute name), and CM (column name).

Connected Hosts attribute

The number of hosts that are connected to the data store. The type is 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.

The following names are defined for this attribute: CONNECTED_HOSTS or CH (warehouse name), Connected Hosts (caption), Connected_Hosts (attribute name), and CH (column name).

Connected VMs attribute

The number of virtual machines that are connected to the data store. The type is 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.

The following names are defined for this attribute: CONNECTED_VMS or CV (warehouse name), Connected VMs (caption), Connected_VMs (attribute name), and CV (column name).

Datacenter attribute

The name of the data center that contains the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Datastore attribute
The name of the data store. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATASTORE (warehouse name), Datastore (caption), Datastore (attribute name), and DATASTORE (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Percent_Used < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Managed System Name attribute
The managed system name of the storage monitoring agent that is associated with the data. The type is 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.

The following names are defined for this attribute: MSN (warehouse name), Managed System Name (caption), MSN (attribute name), and MSN (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Overall Status attribute
The overall status for this data store. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Percent Overcommitted attribute
The percentage of the total capacity of the datastore that is overcommitted.This attribute has a lower bound of -100% and no upper bound. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Unavailable (-21474836). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PERCENT_OVERCOMMITTED or PO (warehouse name), Percent Overcommitted (caption), Percent_Overcommitted (attribute name), and PO (column name).

Percent Used attribute
The percentage of used space in the data store. The type is 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.

The following names are defined for this attribute: PERCENT_USED or PU (warehouse name), Percent Used (caption), Percent_Used (attribute name), and PU (column name).

Remote Host Address attribute
The remote host address for the data store. The type is 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.

The following names are defined for this attribute: REMOTE_HOST_ADDRESS or RHA (warehouse name), Remote Host Address (caption), Remote_Host_Address (attribute name), and RHA (column name).

Remote Path attribute
The remote path for the data store. The type is 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.

The following names are defined for this attribute: REMOTE_PATH or RP (warehouse name), Remote Path (caption), Remote_Path (attribute name), and RP (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Type attribute
The type for the data store. The type is 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.

The following names are defined for this attribute: TYPE (warehouse name), Type (caption), Type (attribute name), and TYPE (column name).

Clustered Resource Pools attribute group

This attribute group describes the resource pools that are members of a cluster. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Cluster Name attribute
The name of the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: CLUSTER_NAME or CN (warehouse name), Cluster Name (caption), Cluster_Name (attribute name), and CN (column name).

CPU Usage attribute
The CPU usage in MHz of all running child virtual machines including virtual machines in child resource pools. The type is 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.

The following names are defined for this attribute: CPU_USAGE (warehouse name), CPU Usage (caption), CPU_Usage (attribute name), and CPU_USAGE (column name).

DataCenter attribute
The name of the data center that contains the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), DataCenter (caption), DataCenter (attribute name), and DATACENTER (column name).

Include Data In Summarization 0 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Max_CPU_Usage < 0) || (Percent_CPU_Usage < 0) || (CPU_Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Max_Memory_Usage < 0) || (Memory_Usage < 0) || (Percent_Memory_Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Max CPU Usage attribute

The current upper bound on CPU usage in MHz.This limit is based on the limit that is configured for the resource pool and the limits that are configured for all parent resource pools. The type is 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.

The following names are defined for this attribute: MAX_CPU_USAGE or MCU (warehouse name), Max CPU Usage (caption), Max_CPU_Usage (attribute name), and MCU (column name).

Max Memory Usage attribute

The current upper bound on memory usage in MB. This limit is based on the limit configured for this resource pool and the limits configured for all parent resource pools. The type is 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.

The following names are defined for this attribute: MAX_MEMORY_USAGE or MMU (warehouse name), Max Memory Usage (caption), Max_Memory_Usage (attribute name), and MMU (column name).

Memory Usage attribute

The memory usage in MB of all running child virtual machines including virtual machines in child resource pools. The type is 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.

The following names are defined for this attribute: MEMORY_USAGE or MU (warehouse name), Memory Usage (caption), Memory_Usage (attribute name), and MU (column name).

Node attribute

The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute

This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

NodeType attribute
The type of node. The valid values are kvm.Resource_Pool and kvm.Virtual_App. The type is string.

The following names are defined for this attribute: NODETYPE (warehouse name), NodeType (caption), NodeType (attribute name), and NODETYPE (column name).

Overall Status attribute
The overall status for this alarm. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Percent CPU Usage attribute
The percentage of CPU resources being used relative to the maximum amount currently available to this resource pool. The type is 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.

The following names are defined for this attribute: PERCENT_CPU_USAGE or PCU (warehouse name), Percent CPU Usage (caption), Percent_CPU_Usage (attribute name), and PCU (column name).

Percent Memory Usage attribute
The percentage of memory resources being used relative to the maximum amount currently available to this resource pool. The type is 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.

The following names are defined for this attribute: PERCENT_MEMORY_USAGE or PMU (warehouse name), Percent Memory Usage (caption), Percent_Memory_Usage (attribute name), and PMU (column name).

Pool Name attribute
The name of this resource pool. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: POOL_NAME (warehouse name), Pool Name (caption), Pool_Name (attribute name), and POOL_NAME (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Clustered Servers attribute group

This attribute group describes the ESX servers that are members of a cluster. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Cluster Name attribute
The name of the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: CLUSTER_NAME or CN (warehouse name), Cluster Name (caption), Cluster_Name (attribute name), and CN (column name).

CPU Effective Contribution attribute

The percentage of CPU resources that this server contributes to the effective CPU of the cluster. The type is 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.

The following names are defined for this attribute: CPU_EFFECTIVE_CONTRIBUTION or CEC (warehouse name), CPU Effective Contribution (caption), CPU_Effective_Contribution (attribute name), and CEC (column name).

CPU Effective Utilization attribute

The CPU usage of this server as a percentage of the effective CPU resources that are owned by this cluster. The type is 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.

The following names are defined for this attribute: CPU_EFFECTIVE_UTILIZATION or CEU (warehouse name), CPU Effective Utilization (caption), CPU_Effective_Utilization (attribute name), and CEU (column name).

CPU Total Contribution attribute

The percentage of CPU resources that this server contributes to the total CPU of the cluster. The type is 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.

The following names are defined for this attribute: CPU_TOTAL_CONTRIBUTION or CTC (warehouse name), CPU Total Contribution (caption), CPU_Total_Contribution (attribute name), and CTC (column name).

CPU Total Utilization attribute

The CPU usage of this server as a percentage of the total CPU resources that are owned by this cluster. The type is 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.

The following names are defined for this attribute: CPU_TOTAL_UTILIZATION or CTU (warehouse name), CPU Total Utilization (caption), CPU_Total_Utilization (attribute name), and CTU (column name).

DataCenter attribute

The name of the data center that contains the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), DataCenter (caption), DataCenter (attribute name), and DATACENTER (column name).

Include Data In Summarization 0 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Server_CPU_Utilization < 0) || (CPU_Total_Utilization < 0) || (CPU_Effective_Utilization < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Server_Memory_Utilization < 0) || (Memory_Total_Utilization < 0) || (Memory_Effective_Utilization < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (CPU_Effective_Contribution < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Include Data In Summarization 3 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Mem_Effective_Contribution < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_3 or IDIS3 (warehouse name), Include Data In Summarization 3 (caption), Include_Data_In_Summarization_3 (attribute name), and IDIS3 (column name).

Mem Effective Contribution attribute

The percentage of memory resources that this server contributes to the effective memory of the cluster. The type is 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.

The following names are defined for this attribute: MEM_EFFECTIVE_CONTRIBUTION or MEC (warehouse name), Mem Effective Contribution (caption), Mem_Effective_Contribution (attribute name), and MEC (column name).

Mem Total Contribution attribute

The percentage of memory resources that this server contributes to the total memory of the cluster. The type is 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.

The following names are defined for this attribute: MEM_TOTAL_CONTRIBUTION or MTC (warehouse name), Mem Total Contribution (caption), Mem_Total_Contribution (attribute name), and MTC (column name).

Memory Effective Utilization attribute

The memory usage of this server as a percentage of the effective memory resources that are owned by this cluster. The type is 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.

The following names are defined for this attribute: MEMORY_EFFECTIVE_UTILIZATION or MEU (warehouse name), Memory Effective Utilization (caption), Memory_Effective_Utilization (attribute name), and MEU (column name).

Memory Total Utilization attribute

The memory usage of this server as a percentage of the total memory resources that are owned by this cluster. The type is 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.

The following names are defined for this attribute: MEMORY_TOTAL_UTILIZATION or MTU (warehouse name), Memory Total Utilization (caption), Memory_Total_Utilization (attribute name), and MTU (column name).

MSN Name attribute
The managed system name that is associated with the data. The type is 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.

The following names are defined for this attribute: MSN_NAME (warehouse name), MSN Name (caption), MSN_Name (attribute name), and MSN_NAME (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Overall Status attribute
The overall status for this alarm. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Server CPU Utilization attribute
The overall CPU usage of this ESX server. The type is 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.

The following names are defined for this attribute: SERVER_CPU_UTILIZATION or SCU (warehouse name), Server CPU Utilization (caption), Server_CPU_Utilization (attribute name), and SCU (column name).

Server Hostname attribute
The host name of the ESX server that is a member of this cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Server Memory Utilization attribute
The overall memory usage of this ESX server. The type is 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.

The following names are defined for this attribute: SERVER_MEMORY_UTILIZATION or SMU (warehouse name), Server Memory Utilization (caption), Server_Memory_Utilization (attribute name), and SMU (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Clustered Virtual Apps attribute group

This attribute group provides information about the virtual machines and virtual applications in the cluster. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Cluster Name attribute
The name of the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: CLUSTER_NAME or CN (warehouse name), Cluster Name (caption), Cluster_Name (attribute name), and CN (column name).

Datacenter attribute
The name of the data center that contains the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Destroy With Parent attribute
Indicates whether the virtual machine must be removed when the virtual application that is associated with the virtual machine is removed. The type is integer with enumerated values. The following values are defined: Yes (1), No (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DESTROY_WITH_PARENT or DWP (warehouse name), Destroy With Parent (caption), Destroy_With_Parent (attribute name), and DWP (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Start Action attribute
Indicates the method by which the virtual machine starts.The valid values are none and powerOn. The type is 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.

The following names are defined for this attribute: START_ACTION or SA (warehouse name), Start Action (caption), Start_Action (attribute name), and SA (column name).

Start Delay attribute
The amount of time (in seconds) that the subsequent virtual machine was delayed to start in a sequence of virtual machines. The type is 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.

The following names are defined for this attribute: START_DELAY or SD (warehouse name), Start Delay (caption), Start_Delay (attribute name), and SD (column name).

Start Order attribute
Indicates the order in which the virtual machine starts. The type is 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.

The following names are defined for this attribute: START_ORDER or SO (warehouse name), Start Order (caption), Start_Order (attribute name), and SO (column name).

Stop Action attribute
Indicates the method by which the virtual machine stops.The valid values are none, powerOff, guestShutdown, and suspend. The type is 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.

The following names are defined for this attribute: STOP_ACTION or SA0 (warehouse name), Stop Action (caption), Stop_Action (attribute name), and SA0 (column name).

Stop Delay attribute
The amount of time (in seconds) that the subsequent virtual machine was delayed to stop in a sequence of virtual machines. The type is 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.

The following names are defined for this attribute: STOP_DELAY (warehouse name), Stop Delay (caption), Stop_Delay (attribute name), and STOP_DELAY (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Virtual App Name attribute
The name of the virtual application. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VIRTUAL_APP_NAME or VAN (warehouse name), Virtual App Name (caption), Virtual_App_Name (attribute name), and VAN (column name).

Virtual Machine Name attribute
The name of the virtual machine. The type is 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.

The following names are defined for this attribute: VIRTUAL_MACHINE_NAME or VMN (warehouse name), Virtual Machine Name (caption), Virtual_Machine_Name (attribute name), and VMN (column name).

VM MORef attribute
The internal managed object reference name of the virtual machine. The type is 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.

The following names are defined for this attribute: MOREF (warehouse name), VM MORef (caption), MORef (attribute name), and MOREF (column name).

Waiting for Guest attribute
Indicates whether the virtual machine must start after receiving a heartbeat from the guest operating system. The type is integer with enumerated values. The following values are defined: Yes (1), No (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: WAITING_FOR_GUEST or WFG (warehouse name), Waiting for Guest (caption), Waiting_For_Guest (attribute name), and WFG (column name).

Clustered Virtual Machines attribute group

This attribute group describes the virtual machines that are members of a cluster. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Cluster MORef attribute
The internal managed object reference name of this cluster. The type is 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.

The following names are defined for this attribute: CLUSTER_MOREF or CM (warehouse name), Cluster MORef (caption), Cluster_MORef (attribute name), and CM (column name).

Cluster Name attribute
The name of the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: CLUSTER_NAME or CN (warehouse name), Cluster Name (caption), Cluster_Name (attribute name), and CN (column name).

CPU Utilization attribute
The overall CPU usage of this virtual machine during the collection interval. The type is 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.

The following names are defined for this attribute: CPU_UTILIZATION or CU (warehouse name), CPU Utilization (caption), CPU_Utilization (attribute name), and CU (column name).

DataCenter attribute
The name of the data center that contains the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), DataCenter (caption), DataCenter (attribute name), and DATACENTER (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (CPU_Utilization < 0) || (Memory_Utilization < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Memory Utilization attribute
The overall memory usage of this virtual machine during the collection interval. The type is 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.

The following names are defined for this attribute: MEMORY_UTILIZATION or MU (warehouse name), Memory Utilization (caption), Memory_Utilization (attribute name), and MU (column name).

MSN Name attribute
The managed system name that is associated with the data. The type is 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.

The following names are defined for this attribute: MSN_NAME (warehouse name), MSN Name (caption), MSN_Name (attribute name), and MSN_NAME (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Overall Status attribute
The overall status for this alarm. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

VM Name attribute
The user-defined display name of this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_NAME (warehouse name), VM Name (caption), VM_Name (attribute name), and VM_NAME (column name).

Clusters attribute group

This attribute group contains metrics that describe the configuration and performance of a cluster. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Cluster MORef attribute
The internal managed object reference name of this cluster. The type is 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.

The following names are defined for this attribute: CLUSTER_MOREF or CM (warehouse name), Cluster MORef (caption), Cluster_MORef (attribute name), and CM (column name).

Cluster Name attribute
The name of the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: CLUSTER_NAME or CN (warehouse name), Cluster Name (caption), Cluster_Name (attribute name), and CN (column name).

CPU 00 10 attribute
The number of servers in this cluster whose CPU usage is 0 - 10 percent. The type is 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.

The following names are defined for this attribute: CPU_00_10 (warehouse name), CPU 00 10 (caption), CPU_00_10 (attribute name), and CPU_00_10 (column name).

CPU 10 20 attribute
The number of servers in this cluster whose CPU usage is 11 - 20 percent. The type is 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.

The following names are defined for this attribute: CPU_10_20 (warehouse name), CPU 10 20 (caption), CPU_10_20 (attribute name), and CPU_10_20 (column name).

CPU 20 30 attribute
The number of servers in this cluster whose CPU usage is 21 - 30 percent. The type is 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.

The following names are defined for this attribute: CPU_20_30 (warehouse name), CPU 20 30 (caption), CPU_20_30 (attribute name), and CPU_20_30 (column name).

CPU 30 40 attribute
The number of servers in this cluster whose CPU usage is 31 - 40 percent. The type is 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.

The following names are defined for this attribute: CPU_30_40 (warehouse name), CPU 30 40 (caption), CPU_30_40 (attribute name), and CPU_30_40 (column name).

CPU 40 50 attribute
The number of servers in this cluster whose CPU usage is 41 - 50 percent. The type is 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.

The following names are defined for this attribute: CPU_40_50 (warehouse name), CPU 40 50 (caption), CPU_40_50 (attribute name), and CPU_40_50 (column name).

CPU 50 60 attribute
The number of servers in this cluster whose CPU usage is 51 - 60 percent. The type is 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.

The following names are defined for this attribute: CPU_50_60 (warehouse name), CPU 50 60 (caption), CPU_50_60 (attribute name), and CPU_50_60 (column name).

CPU 60 70 attribute
The number of servers in this cluster whose CPU usage is 61 - 70 percent. The type is 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.

The following names are defined for this attribute: CPU_60_70 (warehouse name), CPU 60 70 (caption), CPU_60_70 (attribute name), and CPU_60_70 (column name).

CPU 70 80 attribute
The number of servers in this cluster whose CPU usage is 71 - 80 percent. The type is 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.

The following names are defined for this attribute: CPU_70_80 (warehouse name), CPU 70 80 (caption), CPU_70_80 (attribute name), and CPU_70_80 (column name).

CPU 80 90 attribute
The number of servers in this cluster whose CPU usage is 81 - 90 percent. The type is 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.

The following names are defined for this attribute: CPU_80_90 (warehouse name), CPU 80 90 (caption), CPU_80_90 (attribute name), and CPU_80_90 (column name).

CPU 90 100 attribute
The number of servers in this cluster whose CPU usage is 91 - 100 percent. The type is 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.

The following names are defined for this attribute: CPU_90_100 (warehouse name), CPU 90 100 (caption), CPU_90_100 (attribute name), and CPU_90_100 (column name).

CPU Utilization attribute
The total number of CPU resources being used by the member servers divided by the total CPU of the cluster, excluding any members in maintenance mode. The type is real number (32-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.

The following names are defined for this attribute: CPU_UTILIZATION or CU (warehouse name), CPU Utilization (caption), CPU_Utilization (attribute name), and CU (column name).

Current EVC Mode attribute
The current Enhanced VMotion Compatibility (EVC) mode of the cluster. The type is 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.

The following names are defined for this attribute: CURRENT_EVC_MODE or CEM (warehouse name), Current EVC Mode (caption), Current_EVC_Mode (attribute name), and CEM (column name).

DataCenter attribute
The name of the data center that contains the cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), DataCenter (caption), DataCenter (attribute name), and DATACENTER (column name).

Datacenter MORef attribute
The internal managed object reference name of the data center for this cluster. The type is 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.

The following names are defined for this attribute: DATACENTER_MOREF or DM (warehouse name), Datacenter MORef (caption), Datacenter_MORef (attribute name), and DM (column name).

Datastores Total Free Space attribute
The total free space of all data stores connected to this cluster in MB. The type is 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.

The following names are defined for this attribute: DATASTORES_TOTAL_FREE_SPACE or DTFS (warehouse name), Datastores Total Free Space (caption), Datastores_Total_Free_Space (attribute name), and DTFS (column name).

Datastores Total Space attribute
The total space of all data stores connected to this cluster in MB. The type is 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.

The following names are defined for this attribute: DATASTORES_TOTAL_SPACE or DTS (warehouse name), Datastores Total Space (caption), Datastores_Total_Space (attribute name), and DTS (column name).

DRS Enabled attribute

Indicates whether the VMware Dynamic Resource Scheduling facility is enabled for this cluster. The type is integer with enumerated values. The following values are defined: Yes (1), No (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DRS_ENABLED or DE (warehouse name), DRS Enabled (caption), DRS_Enabled (attribute name), and DE (column name).

Effective CPU attribute

The amount of CPU in GHz that is available to run virtual machines.This is an aggregation from all servers that are running normally.The amount of CPU used by the service consoles on each server is not included in the total. The type is real number (32-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.

The following names are defined for this attribute: EFFECTIVE_CPU or EC (warehouse name), Effective CPU (caption), Effective_CPU (attribute name), and EC (column name).

Effective Memory attribute

The amount of memory in GB that is available to run virtual machines. The type is real number (32-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.

The following names are defined for this attribute: EFFECTIVE_MEMORY or EM (warehouse name), Effective Memory (caption), Effective_Memory (attribute name), and EM (column name).

Effective Servers attribute

The number of ESX servers that are available to run virtual machines.Hosts that are unresponsive or in VMware maintenance mode are not counted. The type is 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.

The following names are defined for this attribute: EFFECTIVE_SERVERS or ES (warehouse name), Effective Servers (caption), Effective_Servers (attribute name), and ES (column name).

HA Enabled attribute

Indicates whether the VMware High Availability feature is enabled for this cluster. The type is integer with enumerated values. The following values are defined: Yes (1), No (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: HA_ENABLED (warehouse name), HA Enabled (caption), HA_Enabled (attribute name), and HA_ENABLED (column name).

Include Data In Summarization 0 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Number_Servers < 0) || (Number_CPUs < 0) || (Effective_Servers < 0) || (Percent_Effective_Servers < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_Memory < 0) || (Effective_Memory < 0) || (Percent_Effective_Memory < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 10 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_VM_Provisioned_Space < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_10 or IDIS10 (warehouse name), Include Data In Summarization 10 (caption), Include_Data_In_Summarization_10 (attribute name), and IDIS10 (column name).

Include Data In Summarization 2 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_CPU < 0) || (Effective_CPU < 0) || (Percent_Effective_CPU < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Include Data In Summarization 3 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (CPU_Utilization < 0) || (Memory_Utilization < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_3 or IDIS3 (warehouse name), Include Data In Summarization 3 (caption), Include_Data_In_Summarization_3 (attribute name), and IDIS3 (column name).

Include Data In Summarization 4 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (CPU_00_10 < 0) || (CPU_10_20 < 0) || (CPU_20_30 < 0) || (CPU_30_40 < 0) || (CPU_40_50 < 0) || (CPU_50_60 < 0) || (CPU_60_70 < 0) || (CPU_70_80 < 0) || (CPU_80_90 < 0) || (CPU_90_100 < 0) || (Memory_00_10 < 0) || (Memory_10_20 < 0) || (Memory_20_30 < 0) || (Memory_30_40 < 0) || (Memory_40_50 < 0) || (Memory_50_60 < 0) || (Memory_60_70 < 0) || (Memory_70_80 < 0) || (Memory_80_90 < 0) || (Memory_90_100 < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_4 or IDIS4 (warehouse name), Include Data In Summarization 4 (caption), Include_Data_In_Summarization_4 (attribute name), and IDIS4 (column name).

Include Data In Summarization 5 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Number_vMotions < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_5 or IDIS5 (warehouse name), Include Data In Summarization 5 (caption), Include_Data_In_Summarization_5 (attribute name), and IDIS5 (column name).

Include Data In Summarization 6 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Datastores_Total_Free_Space < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_6 or IDIS6 (warehouse name), Include Data In Summarization 6 (caption), Include_Data_In_Summarization_6 (attribute name), and IDIS6 (column name).

Include Data In Summarization 7 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Physical_NICs < 0) || (Physical_NICs_Down < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_7 or IDIS7 (warehouse name), Include Data In Summarization 7 (caption), Include_Data_In_Summarization_7 (attribute name), and IDIS7 (column name).

Include Data In Summarization 8 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Servers_In_Maintenance_Mode < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_8 or IDIS8 (warehouse name), Include Data In Summarization 8 (caption), Include_Data_In_Summarization_8 (attribute name), and IDIS8 (column name).

Include Data In Summarization 9 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_VM_Configured_Memory < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_9 or IDIS9 (warehouse name), Include Data In Summarization 9 (caption), Include_Data_In_Summarization_9 (attribute name), and IDIS9 (column name).

Memory 00 10 attribute

The number of servers in this cluster whose memory usage is 0 - 10 percent. The type is 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.

The following names are defined for this attribute: MEMORY_00_10 or M01 (warehouse name), Memory 00 10 (caption), Memory_00_10 (attribute name), and M01 (column name).

Memory 10 20 attribute

The number of servers in this cluster whose memory usage is 11 - 20 percent. The type is 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.

The following names are defined for this attribute: MEMORY_10_20 or M12 (warehouse name), Memory 10 20 (caption), Memory_10_20 (attribute name), and M12 (column name).

Memory 20 30 attribute

The number of servers in this cluster whose memory usage is 21 - 30 percent. The type is 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.

The following names are defined for this attribute: MEMORY_20_30 or M23 (warehouse name), Memory 20 30 (caption), Memory_20_30 (attribute name), and M23 (column name).

Memory 30 40 attribute
The number of servers in this cluster whose memory usage is 31 - 40 percent. The type is 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.

The following names are defined for this attribute: MEMORY_30_40 or M34 (warehouse name), Memory 30 40 (caption), Memory_30_40 (attribute name), and M34 (column name).

Memory 40 50 attribute
The number of servers in this cluster whose memory usage is 41 - 50 percent. The type is 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.

The following names are defined for this attribute: MEMORY_40_50 or M45 (warehouse name), Memory 40 50 (caption), Memory_40_50 (attribute name), and M45 (column name).

Memory 50 60 attribute
The number of servers in this cluster whose memory usage is 51 - 60 percent. The type is 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.

The following names are defined for this attribute: MEMORY_50_60 or M56 (warehouse name), Memory 50 60 (caption), Memory_50_60 (attribute name), and M56 (column name).

Memory 60 70 attribute
The number of servers in this cluster whose memory usage is 61 - 70 percent. The type is 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.

The following names are defined for this attribute: MEMORY_60_70 or M67 (warehouse name), Memory 60 70 (caption), Memory_60_70 (attribute name), and M67 (column name).

Memory 70 80 attribute
The number of servers in this cluster whose memory usage is 71 - 80 percent. The type is 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.

The following names are defined for this attribute: MEMORY_70_80 or M78 (warehouse name), Memory 70 80 (caption), Memory_70_80 (attribute name), and M78 (column name).

Memory 80 90 attribute
The number of servers in this cluster whose memory usage is 81 - 90 percent. The type is 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.

The following names are defined for this attribute: MEMORY_80_90 or M89 (warehouse name), Memory 80 90 (caption), Memory_80_90 (attribute name), and M89 (column name).

Memory 90 100 attribute
The number of servers in this cluster whose memory usage is 91 - 100 percent. The type is 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.

The following names are defined for this attribute: MEMORY_90_100 or M91 (warehouse name), Memory 90 100 (caption), Memory_90_100 (attribute name), and M91 (column name).

Memory Utilization attribute
The total amount of memory resources being used by the member servers divided by the total memory of the cluster, excluding any members in maintenance mode. The type is real number (32-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.

The following names are defined for this attribute: MEMORY_UTILIZATION or MU (warehouse name), Memory Utilization (caption), Memory_Utilization (attribute name), and MU (column name).

Node attribute

The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute

This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Number CPUs attribute

The number of physical CPU cores across the cluster. The type is 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.

The following names are defined for this attribute: NUMBER_CPUS or NC (warehouse name), Number CPUs (caption), Number_CPUs (attribute name), and NC (column name).

Number Servers attribute

The number of ESX servers that are members of this cluster. The type is 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.

The following names are defined for this attribute: NUMBER_SERVERS or NS (warehouse name), Number Servers (caption), Number_Servers (attribute name), and NS (column name).

Number vMotions attribute

The total number of migrations that have occurred within this cluster. The type is 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.

The following names are defined for this attribute: NUMBER_VMOTIONS or NV (warehouse name), Number vMotions (caption), Number_vMotions (attribute name), and NV (column name).

Number VMs attribute

The number of virtual machines that are configured within this cluster. The type is 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.

The following names are defined for this attribute: NUMBER_VMS (warehouse name), Number VMs (caption), Number_VMs (attribute name), and NUMBER_VMS (column name).

Number VMs On attribute

The number of virtual machines that are configured within this cluster that are powered on. The type is 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.

The following names are defined for this attribute: NUMBER_VMS_ON or NVO (warehouse name), Number VMs On (caption), Number_VMs_On (attribute name), and NVO (column name).

Overall Status attribute

The overall operational status of the cluster. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Percent Datastore Usage attribute
The percentage of datastore used by the cluster. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: PERCENT_DATASTORE_USAGE or PDU (warehouse name), Percent Datastore Usage (caption), Percent_Datastore_Usage (attribute name), and PDU (column name).

Percent Effective CPU attribute
The percentage of CPU for the cluster that is available to run virtual machines. The type is 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.

The following names are defined for this attribute: PERCENT_EFFECTIVE_CPU or PEC (warehouse name), Percent Effective CPU (caption), Percent_Effective_CPU (attribute name), and PEC (column name).

Percent Effective Memory attribute
The percentage of memory for the cluster that is available to run virtual machines. The type is 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.

The following names are defined for this attribute: PERCENT_EFFECTIVE_MEMORY or PEM (warehouse name), Percent Effective Memory (caption), Percent_Effective_Memory (attribute name), and PEM (column name).

Percent Effective Servers attribute
The percentage of servers defined to the cluster that are available to run virtual machines. The type is 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.

The following names are defined for this attribute: PERCENT_EFFECTIVE_SERVERS or PES (warehouse name), Percent Effective Servers (caption), Percent_Effective_Servers (attribute name), and PES (column name).

Physical NICs attribute
The total number of physical network interface cards in the cluster. The type is 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.

The following names are defined for this attribute: PHYSICAL_NICS or PN (warehouse name), Physical NICs (caption), Physical_NICs (attribute name), and PN (column name).

Physical NICs Down attribute
The total number of physical network interface cards in the cluster with a link status of down. The type is 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.

The following names are defined for this attribute: PHYSICAL_NICS_DOWN or PND (warehouse name), Physical NICs Down (caption), Physical_NICs_Down (attribute name), and PND (column name).

Servers In Maintenance Mode attribute
The number of ESX servers that are in maintenance mode. The type is 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.

The following names are defined for this attribute: SERVERS_IN_MAINTENANCE_MODE or SIMM (warehouse name), Servers In Maintenance Mode (caption), Servers_In_Maintenance_Mode (attribute name), and SIMM (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total CPU attribute
The total amount of CPU resources in GHz over all of the member servers in the cluster. The type is real number (32-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.

The following names are defined for this attribute: TOTAL_CPU (warehouse name), Total CPU (caption), Total_CPU (attribute name), and TOTAL_CPU (column name).

Total Memory attribute
The total memory capacity in GB over all of the member servers in the cluster. The type is real number (32-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.

The following names are defined for this attribute: TOTAL_MEMORY or TM (warehouse name), Total Memory (caption), Total_Memory (attribute name), and TM (column name).

Total VM Configured Memory attribute
The total amount of memory in GB configured for all VMs in the cluster. The type is real number (32-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.

The following names are defined for this attribute: TOTAL_VM_CONFIGURED_MEMORY or TVCM (warehouse name), Total VM Configured Memory (caption), Total_VM_Configured_Memory (attribute name), and TVCM (column name).

Total VM Provisioned Space attribute
The total amount of space in GB that is provisioned for use by VMs in this cluster. The type is real number (32-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.

The following names are defined for this attribute: TOTAL_VM_PROVISIONED_SPACE or TVPS (warehouse name), Total VM Provisioned Space (caption), Total_VM_Provisioned_Space (attribute name), and TVPS (column name).

Datacenters attribute group

This attribute group contains information about the data centers in the virtual infrastructure. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

CPU Utilization attribute
The percent of available CPU being used in this data center. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: CPU_UTILIZATION or CU (warehouse name), CPU Utilization (caption), CPU_Utilization (attribute name), and CU (column name).

Datacenter attribute
The name of this data center. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Effective CPU attribute
The total amount of effective CPU of this data center in MHz. The type is 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.

The following names are defined for this attribute: EFFECTIVE_CPU or EC (warehouse name), Effective CPU (caption), Effective_CPU (attribute name), and EC (column name).

Effective Memory attribute
The total amount of effective memory of this data center in MB. The type is 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.

The following names are defined for this attribute: EFFECTIVE_MEMORY or EM (warehouse name), Effective Memory (caption), Effective_Memory (attribute name), and EM (column name).

Effective Servers attribute
The total number of effective servers that are members of this data center. The type is 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.

The following names are defined for this attribute: EFFECTIVE_SERVERS or ES (warehouse name), Effective Servers (caption), Effective_Servers (attribute name), and ES (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_Servers < 0) || (Effective_Servers < 0) || (Percent_Effective_Servers < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_Memory < 0) || (Effective_Memory < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_CPU < 0) || (Effective_CPU < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Include Data In Summarization 3 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Memory_Utilization < 0) || (CPU_Utilization < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_3 or IDIS3 (warehouse name), Include Data In Summarization 3 (caption), Include_Data_In_Summarization_3 (attribute name), and IDIS3 (column name).

Memory Utilization attribute
The percent of available memory being used in this data center. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: MEMORY_UTILIZATION or MU (warehouse name), Memory Utilization (caption), Memory_Utilization (attribute name), and MU (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Overall Status attribute
The overall status for this alarm. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Percent Effective Servers attribute
The percent of servers that are effective for this data center. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: PERCENT_EFFECTIVE_SERVERS or PES (warehouse name), Percent Effective Servers (caption), Percent_Effective_Servers (attribute name), and PES (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total CPU attribute
The total amount of CPU of this data center in MHz. The type is 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.

The following names are defined for this attribute: TOTAL_CPU (warehouse name), Total CPU (caption), Total_CPU (attribute name), and TOTAL_CPU (column name).

Total Memory attribute
The total amount of memory of this data center in MB. The type is 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.

The following names are defined for this attribute: TOTAL_MEMORY or TM (warehouse name), Total Memory (caption), Total_Memory (attribute name), and TM (column name).

Total Servers attribute
The total numbers of servers that are members of this data center. The type is 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.

The following names are defined for this attribute: TOTAL_SERVERS or TS (warehouse name), Total Servers (caption), Total_Servers (attribute name), and TS (column name).

Datastore Cluster attribute group

This attribute group contains attributes that provide information about the data store cluster (StoragePod). This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Capacity Used attribute
The amount of allocated storage in GB for the data store cluster. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: CAPACITY_USED or CU (warehouse name), Capacity Used (caption), Capacity_Used (attribute name), and CU (column name).

Config Status attribute
The configuration status of the data store cluster.If a problem is detected in the configuration of the data store cluster, the value is displayed as red; and if a problem is about to occur or a transient condition has occurred, the value is displayed as yellow. The type is 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.

The following names are defined for this attribute: CONFIG_STATUS or CS (warehouse name), Config Status (caption), Config_Status (attribute name), and CS (column name).

DataCenter attribute
The name of the data center that the data store cluster belongs to. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), DataCenter (caption), DataCenter (attribute name), and DATACENTER (column name).

Datastore Cluster attribute
The name of the data store cluster. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATASTORE_CLUSTER or DC (warehouse name), Datastore Cluster (caption), Datastore_Cluster (attribute name), and DC (column name).

Datastore Count attribute
The number of data stores in the data store cluster. The type is 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.

The following names are defined for this attribute: DATASTORE_COUNT or DC0 (warehouse name), Datastore Count (caption), Datastore_Count (attribute name), and DC0 (column name).

Default IntraVm Affinity attribute
Indicates whether, by default, each virtual machine must have a virtual disk on the same data store in the data store cluster.The valid values are True and False. The type is integer with enumerated values. The following values are defined: True (1), False (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DEFAULT_INTRAVM_AFFINITY or DIA (warehouse name), Default IntraVm Affinity (caption), Default_IntraVm_Affinity (attribute name), and DIA (column name).

Include Data In Summarization 0 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Capacity < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Percent_Capacity_Free < 0) || (Capacity_Used < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

IO Load Balance Enabled attribute

Indicates whether the data store cluster considers the Storage I/O workload while creating load balancing and initial placement recommendations. The type is integer with enumerated values. The following values are defined: True (1), False (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: IO_LOAD_BALANCE_ENABLED or ILBE (warehouse name), IO Load Balance Enabled (caption), IO_Load_Balance_Enabled (attribute name), and ILBE (column name).

Load Balance Interval attribute

The interval (in minutes) that the Storage Distributed Resource Scheduler (DRS) runs to load balance among data stores within the data store cluster. The type is 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.

The following names are defined for this attribute: LOAD_BALANCE_INTERVAL or LBI (warehouse name), Load Balance Interval (caption), Load_Balance_Interval (attribute name), and LBI (column name).

Node attribute

The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Overall Status attribute

The overall alarm status of the data store cluster.If an alarm is triggered for the data store cluster, the value is displayed as red or yellow. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Percent Capacity Free attribute
The percentage of unused capacity in the data store cluster. The type is 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.

The following names are defined for this attribute: PERCENT_CAPACITY_FREE or PCF (warehouse name), Percent Capacity Free (caption), Percent_Capacity_Free (attribute name), and PCF (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Capacity attribute
The storage capacity in GB of this data store cluster. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: CAPACITY (warehouse name), Total Capacity (caption), Capacity (attribute name), and CAPACITY (column name).

Datastore Host Disks attribute group

This attribute group contains a mapping from a data store to a host disk. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

DataCenter attribute
The name of the data center that contains this disk. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), DataCenter (caption), DataCenter (attribute name), and DATACENTER (column name).

Datastore attribute
The name of the data store on this disk. The type is 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.

The following names are defined for this attribute: DATASTORE (warehouse name), Datastore (caption), Datastore (attribute name), and DATASTORE (column name).

Disk attribute
The name of the disk. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DISK (warehouse name), Disk (caption), Disk (attribute name), and DISK (column name).

Host attribute
The name of the host system for this disk. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: HOST (warehouse name), Host (caption), Host (attribute name), and HOST (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Datastore Topology attribute group

This attribute group contains information about the storage topology of the virtual infrastructure. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

ConnectionType attribute
The connection type from this node to the parent of this node. The type is string.

The following names are defined for this attribute: CONNECTIONTYPE or C0 (warehouse name), ConnectionType (caption), ConnectionType (attribute name), and C0 (column name).

ConnectToNode attribute
Indicates a connection from the NodeID to the node specified here. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: CONNECTTONODE or C (warehouse name), ConnectToNode (caption), ConnectToNode (attribute name), and C (column name).

Datacenter attribute
The name of this data center. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Managed System Name attribute
The managed system name that is associated with the data. The type is 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.

The following names are defined for this attribute: MSN (warehouse name), Managed System Name (caption), MSN (attribute name), and MSN (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
The identifier for this node in the topology. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

NodeName attribute
The name of this node in the tree. The type is string.

The following names are defined for this attribute: NODENAME (warehouse name), NodeName (caption), NodeName (attribute name), and NODENAME (column name).

NodeStatus attribute
The status of this node. The type is string.

The following names are defined for this attribute: NODESTATUS (warehouse name), NodeStatus (caption), NodeStatus (attribute name), and NODESTATUS (column name).

NodeType attribute
The kind of node in the tree. The type is string.

The following names are defined for this attribute: NODETYPE (warehouse name), NodeType (caption), NodeType (attribute name), and NODETYPE (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Datastores attribute group

This attribute group displays general information about data stores. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Accessible attribute
Whether the data store is accessible or not. The type is integer with enumerated values. The following values are defined: Yes (1), No (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ACCESSIBLE (warehouse name), Accessible (caption), Accessible (attribute name), and ACCESSIBLE (column name).

Capacity attribute
The storage capacity in MB of the data store. The type is 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.

The following names are defined for this attribute: CAPACITY (warehouse name), Capacity (caption), Capacity (attribute name), and CAPACITY (column name).

Connected Clusters attribute
The number of clusters with hosts connected to this data store. The type is 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.

The following names are defined for this attribute: CONNECTED_CLUSTERS or CC (warehouse name), Connected Clusters (caption), Connected_Clusters (attribute name), and CC (column name).

Connected Hosts attribute
The number of hosts that are connected to the data store. The type is 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.

The following names are defined for this attribute: CONNECTED_HOSTS or CH (warehouse name), Connected Hosts (caption), Connected_Hosts (attribute name), and CH (column name).

Connected VMs attribute
The number of virtual machines that are connected to the data store. The type is 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.

The following names are defined for this attribute: CONNECTED_VMS or CV (warehouse name), Connected VMs (caption), Connected_VMs (attribute name), and CV (column name).

Datacenter attribute
The name of the data center for the data store. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Datastore Cluster attribute
The name of the data store cluster that the data store belongs to. The type is 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.

The following names are defined for this attribute: DATASTORE_CLUSTER or DC (warehouse name), Datastore Cluster (caption), Datastore_Cluster (attribute name), and DC (column name).

Datastore MORef attribute
The internal managed object reference name of the datastore. The type is 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.

The following names are defined for this attribute: DATASTORE_MOREF or DM (warehouse name), Datastore MORef (caption), Datastore_MORef (attribute name), and DM (column name).

Free Space attribute
The amount of available storage in MB for the data store. The type is 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.

The following names are defined for this attribute: FREE_SPACE (warehouse name), Free Space (caption), Free_Space (attribute name), and FREE_SPACE (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Used_Space < 0) || (Free_Space < 0) || (Percent_Used < 0) || (Percent_Free < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Connected_Hosts < 0) || (Connected_VMs < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Include Data In Summarization 3 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_Read_KBps < 0) || (Total_Write_KBps < 0) || (Total_IO_KBps < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_3 or IDIS3 (warehouse name), Include Data In Summarization 3 (caption), Include_Data_In_Summarization_3 (attribute name), and IDIS3 (column name).

Include Data In Summarization 4 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Percent_Overcommitted == -2147483600)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_4 or IDIS4 (warehouse name), Include Data In Summarization 4 (caption), Include_Data_In_Summarization_4 (attribute name), and IDIS4 (column name).

Include Data In Summarization 5 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Snapshot_Storage_Consumed < 0) || (Percent_Snapshot_Storage_Consumed < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_5 or IDIS5 (warehouse name), Include Data In Summarization 5 (caption), Include_Data_In_Summarization_5 (attribute name), and IDIS5 (column name).

Managed System Name attribute
The managed system name of the storage monitoring agent that is associated with the data. The type is 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.

The following names are defined for this attribute: MSN (warehouse name), Managed System Name (caption), MSN (attribute name), and MSN (column name).

Maximum File Size attribute
The maximum size in KB of a file that might be allocated on this data store. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-1), > 2048GB (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAXIMUM_FILE_SIZE or MFS (warehouse name), Maximum File Size (caption), Maximum_File_Size (attribute name), and MFS (column name).

Name attribute
The name of the data store. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NAME (warehouse name), Name (caption), Name (attribute name), and NAME (column name).

NetApp Volume Name attribute
A best effort guess at the corresponding NetApp volume name for the datastore. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), No DNS Record (No_DNS_Record). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: NETAPP_VOLUME_NAME or NVN (warehouse name), NetApp Volume Name (caption), NetApp_Volume_Name (attribute name), and NVN (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Overall Status attribute
The overall status for the data store. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), red (red), yellow (yellow), green (green), gray (gray). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Overcommitted attribute
The amount of space, in megabytes, that the datastore has provisioned without available backing storage.This value can be negative, with a lower bound of negative free space. The type is 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.

The following names are defined for this attribute: OVERCOMMITTED or O (warehouse name), Overcommitted (caption), Overcommitted (attribute name), and O (column name).

Percent Free attribute
The percentage of unused space in this data store. The type is 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.

The following names are defined for this attribute: PERCENT_FREE or PF (warehouse name), Percent Free (caption), Percent_Free (attribute name), and PF (column name).

Percent Overcommitted attribute
The percentage of the total capacity of the datastore that is overcommitted.This attribute has a lower bound of -100% and no upper bound. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: Unavailable (-21474836). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: PERCENT_OVERCOMMITTED or PO (warehouse name), Percent Overcommitted (caption), Percent_Overcommitted (attribute name), and PO (column name).

Percent Snapshot Storage Consumed attribute
The percentage amount of disk space that is used by the snapshots. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: PERCENT_SNAPSHOT_STORAGE_CONSUMED or PSSC (warehouse name), Percent Snapshot Storage Consumed (caption), Percent_Snapshot_Storage_Consumed (attribute name), and PSSC (column name).

Percent Used attribute
The percentage of used space in the data store. The type is 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.

The following names are defined for this attribute: PERCENT_USED or PU (warehouse name), Percent Used (caption), Percent_Used (attribute name), and PU (column name).

Remote Host Address attribute
The remote host address for the data store. The type is 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.

The following names are defined for this attribute: REMOTE_HOST_ADDRESS or RHA (warehouse name), Remote Host Address (caption), Remote_Host_Address (attribute name), and RHA (column name).

Remote Path attribute
The remote path for the data store. The type is 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.

The following names are defined for this attribute: REMOTE_PATH or RP (warehouse name), Remote Path (caption), Remote_Path (attribute name), and RP (column name).

Snapshot Storage Consumed attribute
The amount of disk space (in GB) that is used by the snapshots. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: SNAPSHOT_STORAGE_CONSUMED or SSC (warehouse name), Snapshot Storage Consumed (caption), Snapshot_Storage_Consumed (attribute name), and SSC (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total IO attribute
The sum of total kilobytes read and written per second by all virtual machines that are configured for this datastore. The type is 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.

The following names are defined for this attribute: TOTAL_IO_KBPS or TIK (warehouse name), Total IO (caption), Total_IO_KBps (attribute name), and TIK (column name).

Total Read attribute
The total kilobytes read per second by all virtual machines that are configured for this datastore. The type is 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.

The following names are defined for this attribute: TOTAL_READ_KBPS or TRK (warehouse name), Total Read (caption), Total_Read_KBps (attribute name), and TRK (column name).

Total Write attribute
The total kilobytes written per second by all virtual machines that are configured for this datastore. The type is 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.

The following names are defined for this attribute: TOTAL_WRITE_KBPS or TWK (warehouse name), Total Write (caption), Total_Write_KBps (attribute name), and TWK (column name).

Type attribute
The type for the data store. The type is 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.

The following names are defined for this attribute: TYPE (warehouse name), Type (caption), Type (attribute name), and TYPE (column name).

URL attribute
The remote URL for the data store. The type is 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.

The following names are defined for this attribute: URL (warehouse name), URL (caption), URL (attribute name), and URL (column name).

Used Space attribute
The amount of allocated storage in MB for the data store. The type is 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.

The following names are defined for this attribute: USED_SPACE (warehouse name), Used Space (caption), Used_Space (attribute name), and USED_SPACE (column name).

Director attribute group

This attribute group contains information about the IBM Systems Director configuration. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

DirectorPort attribute
The port number for the IBM Director Server. The type is string.

The source for this attribute is Script data

The following names are defined for this attribute: DIRECTORPORT or D0 (warehouse name), DirectorPort (caption), DirectorPort (attribute name), and D0 (column name).

DirectorServer attribute
The host name of the IBM Director Server. The type is string.

The source for this attribute is Script data

The following names are defined for this attribute: DIRECTORSERVER or D (warehouse name), DirectorServer (caption), DirectorServer (attribute name), and D (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

UseTEPCredential attribute
Use Tivoli Enterprise Portal credentials for IBM Systems Director authentication. The type is string.

The source for this attribute is Script data

The following names are defined for this attribute: USETEPCREDENTIAL or U (warehouse name), UseTEPCredential (caption), UseTEPCredential (attribute name), and U (column name).

Distributed Virtual Portgroups attribute group

This attribute group contains information about the distributed virtual portgroups in the virtual infrastructure. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Blocked attribute
Whether traffic is being blocked for this portgroup. The type is 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.

The following names are defined for this attribute: BLOCKED (warehouse name), Blocked (caption), Blocked (attribute name), and BLOCKED (column name).

Datacenter attribute
The name of the data center that uses this distributed virtual portgroup. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Inbound Shaping Average Bandwidth attribute
The inbound traffic shaping target for average bandwidth. The type is 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.

The following names are defined for this attribute: INBOUND_SHAPING_AVERAGE_BANDWIDTH or ISAB (warehouse name), Inbound Shaping Average Bandwidth (caption), Inbound_Shaping_Average_Bandwidth (attribute name), and ISAB (column name).

Inbound Shaping Burst Size attribute
The inbound traffic shaping target for burst size. The type is 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.

The following names are defined for this attribute: INBOUND_SHAPING_BURST_SIZE or ISBS (warehouse name), Inbound Shaping Burst Size (caption), Inbound_Shaping_Burst_Size (attribute name), and ISBS (column name).

Inbound Shaping Enabled attribute

Whether inbound traffic shaping is enabled for this portgroup. The type is 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.

The following names are defined for this attribute: INBOUND_SHAPING_ENABLED or ISE (warehouse name), Inbound Shaping Enabled (caption), Inbound_Shaping_Enabled (attribute name), and ISE (column name).

Inbound Shaping Peak Bandwidth attribute

The inbound traffic shaping target for peak bandwidth. The type is 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.

The following names are defined for this attribute: INBOUND_SHAPING_PEAK_BANDWIDTH or ISPB (warehouse name), Inbound Shaping Peak Bandwidth (caption), Inbound_Shaping_Peak_Bandwidth (attribute name), and ISPB (column name).

Include Data In Summarization 0 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Inbound_Shaping_Average_Bandwidth < 0) || (Inbound_Shaping_Burst_Size < 0) || (Inbound_Shaping_Peak_Bandwidth < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Outbound_Shaping_Peak_Bandwidth < 0) || (Outbound_Shaping_Average_Bandwidth < 0) || (Outbound_Shaping_Burst_Size < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Node attribute

The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Outbound Shaping Average Bandwidth attribute

The outbound traffic shaping target for average bandwidth. The type is 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.

The following names are defined for this attribute: OUTBOUND_SHAPING_AVERAGE_BANDWIDTH or OSAB (warehouse name), Outbound Shaping Average Bandwidth (caption), Outbound_Shaping_Average_Bandwidth (attribute name), and OSAB (column name).

Outbound Shaping Burst Size attribute

The outbound traffic shaping target for burst size. The type is 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.

The following names are defined for this attribute: OUTBOUND_SHAPING_BURST_SIZE or OSBS (warehouse name), Outbound Shaping Burst Size (caption), Outbound_Shaping_Burst_Size (attribute name), and OSBS (column name).

Outbound Shaping Enabled attribute
Whether outbound traffic shaping is enabled for this portgroup. The type is 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.

The following names are defined for this attribute: OUTBOUND_SHAPING_ENABLED or OSE (warehouse name), Outbound Shaping Enabled (caption), Outbound_Shaping_Enabled (attribute name), and OSE (column name).

Outbound Shaping Peak Bandwidth attribute
The outbound traffic shaping target for peak bandwidth. The type is 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.

The following names are defined for this attribute: OUTBOUND_SHAPING_PEAK_BANDWIDTH or OSPB (warehouse name), Outbound Shaping Peak Bandwidth (caption), Outbound_Shaping_Peak_Bandwidth (attribute name), and OSPB (column name).

Overall Status attribute
The overall alarm status of the portgroup. A value of red or yellow indicates that an alarm has been triggered for the portgroup. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Portgroup attribute
The name of this portgroup. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: PORTGROUP_NAME or PN (warehouse name), Portgroup (caption), Portgroup_Name (attribute name), and PN (column name).

Switch attribute
The name of the distributed virtual switch associated with this portgroup. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SWITCH_NAME or SN (warehouse name), Switch (caption), Switch_Name (attribute name), and SN (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Type attribute
The type of this portgroup. The type is 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.

The following names are defined for this attribute: TYPE (warehouse name), Type (caption), Type (attribute name), and TYPE (column name).

VLAN ID attribute
The VLAN ID used by this portgroup. For portgroups that support ranges of VLANs or multiple VLANs, this value is set to Unavailable. The type is 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.

The following names are defined for this attribute: VLAN_ID (warehouse name), VLAN ID (caption), VLAN_ID (attribute name), and VLAN_ID (column name).

VLAN Type attribute
The type of VLAN used for this portgroup. The type is 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.

The following names are defined for this attribute: VLAN_TYPE (warehouse name), VLAN Type (caption), VLAN_Type (attribute name), and VLAN_TYPE (column name).

Distributed Virtual Switch Health attribute group

This attribute group contains information about the health check of host system for distributed virtual switches. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Datacenter attribute
The name of the data center that uses this distributed virtual switch. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

DVS Teaming Status attribute
The teaming check status of the Distributed Virtual Switch. This teaming check status is available only for the VLAN Health check and the Teaming and Failover Health check. The type is 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.

The following names are defined for this attribute: DVS_TEAMING_STATUS or DTS (warehouse name), DVS Teaming Status (caption), DVS_Teaming_Status (attribute name), and DTS (column name).

Health Check Type attribute
The type of the health check. The type is string with enumerated values. The following values are defined: VLAN Health (com.vmware.vim.VMwareDVSMtuHealthCheckResult), MTU Health (com.vmware.vim.VMwareDVSVlanHealthCheckResult), Teaming and Failover Health (com.vmware.vim.VMwareDVSTeamingHealthCheckResult). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: HEALTH_CHECK_TYPE or HCT (warehouse name), Health Check Type (caption), Health_Check_Type (attribute name), and HCT (column name).

Host attribute
The host name of the ESX server that is connected to the Distributed Virtual Switch. The type is 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.

The following names are defined for this attribute: HOST (warehouse name), Host (caption), Host (attribute name), and HOST (column name).

MTU Mismatch attribute
Indicates whether the Maximum Transmission Unit (MTU) configured in the vSphere Distributed Switch is mismatched with the value configured in the Physical NIC.This MTU mismatch status is available only for MTU Health Check type. The type is 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.

The following names are defined for this attribute: MTU_MISMATCH or MM (warehouse name), MTU Mismatch (caption), MTU_Mismatch (attribute name), and MM (column name).

NIC attribute
The name of the physical network interface card (NIC) that is associated with the uplink. The type is 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.

The following names are defined for this attribute: NIC_NAME (warehouse name), NIC (caption), NIC_Name (attribute name), and NIC_NAME (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Portgroup attribute
The name of the portgroup. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: PORTGROUP_NAME or PN (warehouse name), Portgroup (caption), Portgroup_Name (attribute name), and PN (column name).

Source attribute
The host name of the data source. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SOURCE (warehouse name), Source (caption), Source (attribute name), and SOURCE (column name).

Summary attribute
The health check summary. The type is 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.

The following names are defined for this attribute: SUMMARY (warehouse name), Summary (caption), Summary (attribute name), and SUMMARY (column name).

Switch attribute
The name of the Distributed Virtual Switch. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SWITCH_NAME or SN (warehouse name), Switch (caption), Switch_Name (attribute name), and SN (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Uplink attribute
The name of the uplink that is used by the host to connect to the Distributed Virtual Switch. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: UPLINK_NAME or UN (warehouse name), Uplink (caption), Uplink_Name (attribute name), and UN (column name).

Uplink Key attribute
The uplink key that is used by the host to connect to the Distributed Virtual Switch. The type is 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.

The following names are defined for this attribute: UPLINK_KEY (warehouse name), Uplink Key (caption), Uplink_Key (attribute name), and UPLINK_KEY (column name).

Distributed Virtual Switches attribute group

This attribute group contains information about the distributed virtual switches in the virtual infrastructure. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Datacenter attribute
The name of the data center that uses this distributed virtual switch. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Number_Of_Portgroups < 0) || (Number_Of_Uplinks < 0) || (Number_Of_Ports < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Number_Of_Hosts < 0) || (Number_Of_VMs < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Transmitted < 0) || (Received < 0) || (Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Include Data In Summarization 3 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Max_Number_Of_Ports < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_3 or IDIS3 (warehouse name), Include Data In Summarization 3 (caption), Include_Data_In_Summarization_3 (attribute name), and IDIS3 (column name).

Max Number Ports attribute
The maximum number of ports, excluding conflict ports, allowed for this switch. The type is 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.

The following names are defined for this attribute: MAX_NUMBER_OF_PORTS or MNOP (warehouse name), Max Number Ports (caption), Max_Number_Of_Ports (attribute name), and MNOP (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Number Hosts attribute
The number of hosts that are attached to this switch. The type is 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.

The following names are defined for this attribute: NUMBER_OF_HOSTS or NOH (warehouse name), Number Hosts (caption), Number_Of_Hosts (attribute name), and NOH (column name).

Number Of Portgroups attribute
The number of portgroups, including uplink portgroups, attached to this switch. The type is 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.

The following names are defined for this attribute: NUMBER_OF_PORTGROUPS or NOP (warehouse name), Number Of Portgroups (caption), Number_Of_Portgroups (attribute name), and NOP (column name).

Number Ports attribute
The current number of ports, excluding conflict ports, of this switch. The type is 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.

The following names are defined for this attribute: NUMBER_OF_PORTS or NOP0 (warehouse name), Number Ports (caption), Number_Of_Ports (attribute name), and NOP0 (column name).

Number Uplinks attribute
The number of distributed virtual uplinks that are attached to this switch. The type is 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.

The following names are defined for this attribute: NUMBER_OF_UPLINKS or NOU (warehouse name), Number Uplinks (caption), Number_Of_Uplinks (attribute name), and NOU (column name).

Number VMs attribute
The number of virtual machines that are attached to this switch. The type is 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.

The following names are defined for this attribute: NUMBER_OF_VMS or NOV (warehouse name), Number VMs (caption), Number_Of_VMs (attribute name), and NOV (column name).

Overall Status attribute
The overall alarm status of the distributed virtual switch. A value of red or yellow indicates that an alarm has been triggered for the distributed virtual switch. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Received attribute
The total reception rate in KBps of the uplinks on this switch. The type is 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.

The following names are defined for this attribute: RECEIVED (warehouse name), Received (caption), Received (attribute name), and RECEIVED (column name).

Switch attribute
The name of the distributed virtual switch. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SWITCH_NAME or SN (warehouse name), Switch (caption), Switch_Name (attribute name), and SN (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Transmitted attribute
The total transmission rate in KBps of the uplinks on this switch. The type is 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.

The following names are defined for this attribute: TRANSMITTED or T (warehouse name), Transmitted (caption), Transmitted (attribute name), and T (column name).

Usage attribute
The total rate in KBps that the uplinks are transmitting and receiving data on this switch. The type is 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.

The following names are defined for this attribute: USAGE (warehouse name), Usage (caption), Usage (attribute name), and USAGE (column name).

Distributed Virtual Uplinks attribute group

This attribute group contains information about the distributed virtual uplinks in the virtual infrastructure. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Component State attribute
The component state of the uplink. The type is 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.

The following names are defined for this attribute: COMPONENT_STATE or CS (warehouse name), Component State (caption), Component_State (attribute name), and CS (column name).

Datacenter attribute
The name of the data center that uses this distributed virtual uplink. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Duplex attribute
The current operating mode of the NIC. The type is 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.

The following names are defined for this attribute: DUPLEX (warehouse name), Duplex (caption), Duplex (attribute name), and DUPLEX (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Transmitted < 0) || (Received < 0) || (Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Link_Utilization < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Link_Speed < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Link Speed attribute
The current operating speed of the NIC in Mbps. The type is 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.

The following names are defined for this attribute: LINK_SPEED (warehouse name), Link Speed (caption), Link_Speed (attribute name), and LINK_SPEED (column name).

Link Utilization attribute
The percent usage of the NIC relative to the capacity of the link (including duplex). The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: LINK_UTILIZATION or LU (warehouse name), Link Utilization (caption), Link_Utilization (attribute name), and LU (column name).

Managed System Name attribute
The managed system name of the subnode for the ESX server of the uplink. The type is 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.

The following names are defined for this attribute: SUBNODE_MSN or SM (warehouse name), Managed System Name (caption), Subnode_MSN (attribute name), and SM (column name).

NIC attribute
The name of the physical NIC associated with this uplink. The type is 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.

The following names are defined for this attribute: NIC (warehouse name), NIC (caption), NIC (attribute name), and NIC (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Overall Status attribute
The overall alarm status of the uplink. A value of red or yellow indicates that an alarm has been triggered for the uplink. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Portgroup attribute
The name of the portgroup for this uplink. The type is 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.

The following names are defined for this attribute: PORTGROUP_NAME or PN (warehouse name), Portgroup (caption), Portgroup_Name (attribute name), and PN (column name).

Received attribute
The total reception rate in KBps of this uplink's physical NIC. The type is 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.

The following names are defined for this attribute: RECEIVED (warehouse name), Received (caption), Received (attribute name), and RECEIVED (column name).

Server Hostname attribute
The host name of the ESX server to which the uplink belongs. The type is 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.

The following names are defined for this attribute: HOST_SYSTEM or HS (warehouse name), Server Hostname (caption), Host_System (attribute name), and HS (column name).

Status attribute
The current status, up or down, of the NIC. The type is 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.

The following names are defined for this attribute: LINK_STATUS or LS (warehouse name), Status (caption), Link_Status (attribute name), and LS (column name).

Switch attribute
The name of the distributed virtual switch that is attached to this uplink. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SWITCH_NAME or SN (warehouse name), Switch (caption), Switch_Name (attribute name), and SN (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Transmitted attribute
The total transmission rate in KBps of this uplink's physical NIC. The type is 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.

The following names are defined for this attribute: TRANSMITTED or T (warehouse name), Transmitted (caption), Transmitted (attribute name), and T (column name).

Uplink attribute
The name of this uplink. The type is 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.

The following names are defined for this attribute: UPLINK_NAME or UN (warehouse name), Uplink (caption), Uplink_Name (attribute name), and UN (column name).

Usage attribute
The total rate in KBps that data is being transmitted and received data on this uplink's physical NIC. The type is 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.

The following names are defined for this attribute: USAGE (warehouse name), Usage (caption), Usage (attribute name), and USAGE (column name).

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

This attribute group contains the following attributes:

Average Collection Duration attribute
The average duration of all data collections of this group in seconds. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVERAGE_COLLECTION_DURATION or COLAVGD (warehouse name), Average Collection Duration (caption), Average_Collection_Duration (attribute name), and COLAVGD (column name).

Cache Hit Percent attribute
The percentage of external data requests for this group that were satisfied from the cache. The type is real number (32-bit counter) with two decimal places of precision.

The following names are defined for this attribute: CACHE_HIT_PERCENT or CACHPCT (warehouse name), Cache Hit Percent (caption), Cache_Hit_Percent (attribute name), and CACHPCT (column name).

Cache Hits attribute
The number of times an external data request for this group was satisfied from the cache. The type is integer (32-bit counter).

The following names are defined for this attribute: CACHE_HITS or CACHEHT (warehouse name), Cache Hits (caption), Cache_Hits (attribute name), and CACHEHT (column name).

Cache Misses attribute
The number of times an external data request for this group was not available in the cache. The type is integer (32-bit counter).

The following names are defined for this attribute: CACHE_MISSES or CACHEMS (warehouse name), Cache Misses (caption), Cache_Misses (attribute name), and CACHEMS (column name).

Error Code attribute
The error code associated with the query. The type is 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), KVM NO DATASOURCES (1000), KVM DATASOURCE LOGIN FAILED (1005), KVM DATASOURCE NOT FOUND (1010), SUBNODE UNAVAILABLE (1033), KVM PROVIDER RESET (2222). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ERROR_CODE or ERRCODE (warehouse name), Error Code (caption), Error_Code (attribute name), and ERRCODE (column name).

Intervals Skipped attribute
The number of times a background data collection for this group was skipped because the previous collection was still running when the next one was due to start. The type is integer (32-bit counter).

The following names are defined for this attribute: INTERVALS_SKIPPED or INTSKIP (warehouse name), Intervals Skipped (caption), Intervals_Skipped (attribute name), and INTSKIP (column name).

Last Collection Duration attribute
The duration of the most recently completed data collection of this group in seconds. The type is real number (32-bit counter) with two decimal places of precision.

The following names are defined for this attribute: LAST_COLLECTION_DURATION or COLDURA (warehouse name), Last Collection Duration (caption), Last_Collection_Duration (attribute name), and COLDURA (column name).

Last Collection Finished attribute
The most recent time a data collection of this group finished. The type is timestamp with enumerated values. The following values are defined: NOT COLLECTED (0691231190000000), NOT COLLECTED (0000000000000001). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_COLLECTION_FINISHED or COLFINI (warehouse name), Last Collection Finished (caption), Last_Collection_Finished (attribute name), and COLFINI (column name).

Last Collection Start attribute
The most recent time a data collection of this group started. The type is timestamp with enumerated values. The following values are defined: NOT COLLECTED (0691231190000000), NOT COLLECTED (0000000000000001). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_COLLECTION_START or COLSTRT (warehouse name), Last Collection Start (caption), Last_Collection_Start (attribute name), and COLSTRT (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Number of Collections attribute
The number of data collections for this group since the agent started. The type is integer (32-bit counter).

The following names are defined for this attribute: NUMBER_OF_COLLECTIONS or NUMCOLL (warehouse name), Number of Collections (caption), Number_of_Collections (attribute name), and NUMCOLL (column name).

Object Name attribute
The name of the performance object. The type is string.

The following names are defined for this attribute: OBJECT_NAME or OBJNAME (warehouse name), Object Name (caption), Object_Name (attribute name), and OBJNAME (column name).

Object Status attribute
The status of the performance object. The type is 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.

The following names are defined for this attribute: OBJECT_STATUS or OBJSTTS (warehouse name), Object Status (caption), Object_Status (attribute name), and OBJSTTS (column name).

Object Type attribute
The type of the performance object. The type is 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.

The following names are defined for this attribute: OBJECT_TYPE or OBJTYPE (warehouse name), Object Type (caption), Object_Type (attribute name), and OBJTYPE (column name).

Query Name attribute
The name of the attribute group. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: QUERY_NAME or ATTRGRP (warehouse name), Query Name (caption), Query_Name (attribute name), and ATTRGRP (column name).

Refresh Interval attribute
The interval at which this group is refreshed in seconds. The type is integer (32-bit counter).

The following names are defined for this attribute: REFRESH_INTERVAL or REFRINT (warehouse name), Refresh Interval (caption), Refresh_Interval (attribute name), and REFRINT (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Events attribute group

This attribute group contains events that are not specific to an ESX server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Category attribute
The severity level that is associated with the event by VMware. The type is 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.

The following names are defined for this attribute: CATEGORY (warehouse name), Category (caption), Category (attribute name), and CATEGORY (column name).

Compute Resource attribute
The compute resource that is associated with this event. The type is 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.

The following names are defined for this attribute: COMPUTE_RESOURCE or CR (warehouse name), Compute Resource (caption), Compute_Resource (attribute name), and CR (column name).

Datacenter attribute
The data center that is associated with this event. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Datastore attribute
The name of the data store that is associated with the event. The type is 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.

The following names are defined for this attribute: DATASTORE (warehouse name), Datastore (caption), Datastore (attribute name), and DATASTORE (column name).

Datastore UUID attribute
The Universal Unique ID of the data store that is associated with the event. The type is 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.

The following names are defined for this attribute: DATASTORE_UUID or DU (warehouse name), Datastore UUID (caption), Datastore_UUID (attribute name), and DU (column name).

Entity Type attribute
The type of entity of the event. The type is 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.

The following names are defined for this attribute: ENTITY_TYPE or ET (warehouse name), Entity Type (caption), Entity_Type (attribute name), and ET (column name).

Event attribute
The event data string. The type is 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.

The following names are defined for this attribute: EVENT (warehouse name), Event (caption), Event (attribute name), and EVENT (column name).

Event Seq Number attribute
A sequence number for this event. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: EVENT_SEQ_NUMBER or ESN (warehouse name), Event Seq Number (caption), Event_Seq_Number (attribute name), and ESN (column name).

Event Text attribute
The full event data string. The type is 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.

The following names are defined for this attribute: EVENT_TEXT (warehouse name), Event Text (caption), Event_Text (attribute name), and EVENT_TEXT (column name).

Event Time attribute
The time that the event occurred. The type is 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.

The following names are defined for this attribute: EVENT_TIME (warehouse name), Event Time (caption), Event_Time (attribute name), and EVENT_TIME (column name).

Event Type attribute
The type of event that is given by VMware. The type is 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.

The following names are defined for this attribute: EVENT_TYPE (warehouse name), Event Type (caption), Event_Type (attribute name), and EVENT_TYPE (column name).

Event Type ID attribute
The type ID of the event that is given by VMware.This is unavailable unless the event is an extended event. The type is 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.

The following names are defined for this attribute: EVENT_TYPE_ID or ETI (warehouse name), Event Type ID (caption), Event_Type_ID (attribute name), and ETI (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Event_Seq_Number < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Source Hostname attribute
The host name of the data source that originated this event. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SOURCE_HOSTNAME or SH (warehouse name), Source Hostname (caption), Source_Hostname (attribute name), and SH (column name).

Timestamp attribute
The time the event was generated. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

UserId attribute
The user ID that caused the event. The type is string with enumerated values. The following values are defined: Not applicable (Not applicable), Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: USERID (warehouse name), UserId (caption), UserId (attribute name), and USERID (column name).

Virtual Machine attribute
The virtual machine that is associated with this event. The type is 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.

The following names are defined for this attribute: VIRTUAL_MACHINE or VM (warehouse name), Virtual Machine (caption), Virtual_Machine (attribute name), and VM (column name).

Virtual Machine UUID attribute
The UUID of the virtual machine that is associated with this event. The type is 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.

The following names are defined for this attribute: VIRTUAL_MACHINE_UUID or VMU (warehouse name), Virtual Machine UUID (caption), Virtual_Machine_UUID (attribute name), and VMU (column name).

Monitored Servers attribute group

This attribute group is the current list of ESX servers that are being monitored. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Subnode Affinity attribute
The affinity of this subnode. The type is string.

The following names are defined for this attribute: SUBNODE_AFFINITY or SN_AFFIN (warehouse name), Subnode Affinity (caption), Subnode_Affinity (attribute name), and SN_AFFIN (column name).

Subnode MSN attribute
The Managed System Name of this subnode. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: SUBNODE_MSN or SN_MSN (warehouse name), Subnode MSN (caption), Subnode_MSN (attribute name), and SN_MSN (column name).

Subnode Resource Name attribute
The resource name of this subnode. The type is string.

The following names are defined for this attribute: SUBNODE_RESOURCE_NAME or SN_RES (warehouse name), Subnode Resource Name (caption), Subnode_Resource_Name (attribute name), and SN_RES (column name).

Subnode Resource Name Enhanced attribute
This is enhanced version of subnode_resource_name attribute with increased limit to show 100 characters of ESX Server name. The type is string.

The following names are defined for this attribute: SUBNODE_RESOURCE_NAME_ENHANCED or SRNE (warehouse name), Subnode Resource Name Enhanced (caption), Subnode_Resource_Name_Enhanced (attribute name), and SRNE (column name).

Subnode Type attribute
The node type of this subnode. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: SUBNODE_TYPE or SN_TYPE (warehouse name), Subnode Type (caption), Subnode_Type (attribute name), and SN_TYPE (column name).

Subnode Version attribute
The version of this subnode. The type is string.

The following names are defined for this attribute: SUBNODE_VERSION or SN_VER (warehouse name), Subnode Version (caption), Subnode_Version (attribute name), and SN_VER (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Networked Servers attribute group

This attribute group lists the hosts on each network. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Datacenter attribute
The data center that this network is on. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Transmitted < 0) || (Received < 0) || (Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Managed System Name attribute
The managed system name of the subnode for the ESX server. The type is 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.

The following names are defined for this attribute: SUBNODE_MSN or SM (warehouse name), Managed System Name (caption), Subnode_MSN (attribute name), and SM (column name).

Network attribute
The name of the network. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NETWORK (warehouse name), Network (caption), Network (attribute name), and NETWORK (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Received attribute
The total reception rate in KBps of the host on this virtual switch. The type is 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.

The following names are defined for this attribute: RECEIVED (warehouse name), Received (caption), Received (attribute name), and RECEIVED (column name).

Server Hostname attribute
The host name of the ESX server that is connected to the network. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Switch attribute
The switch that the network uses. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SWITCH (warehouse name), Switch (caption), Switch (attribute name), and SWITCH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Transmitted attribute
The total transmission rate in KBps of the host on this virtual switch. The type is 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.

The following names are defined for this attribute: TRANSMITTED or T (warehouse name), Transmitted (caption), Transmitted (attribute name), and T (column name).

Usage attribute
The total rate in KBps that the host is transmitting and receiving data on this virtual switch. The type is 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.

The following names are defined for this attribute: USAGE (warehouse name), Usage (caption), Usage (attribute name), and USAGE (column name).

Networked Virtual Machines attribute group

This attribute group lists the virtual machine NICs on each network. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Datacenter attribute
The data center that this virtual machine NIC is on. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Transmitted < 0) || (Received < 0) || (Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Managed System Name attribute
The managed system name of the subnode for the ESX server of the virtual machine NIC. The type is 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.

The following names are defined for this attribute: SUBNODE_MSN or SM (warehouse name), Managed System Name (caption), Subnode_MSN (attribute name), and SM (column name).

Network attribute
The name of the network the virtual machine NIC is on. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NETWORK (warehouse name), Network (caption), Network (attribute name), and NETWORK (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Received attribute
The total reception rate in KBps of this virtual machine NIC. The type is 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.

The following names are defined for this attribute: RECEIVED (warehouse name), Received (caption), Received (attribute name), and RECEIVED (column name).

Server Hostname attribute
The hostname of the ESX server on which the virtual machine resides. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Switch attribute
The name of the virtual switch to which the virtual machine NIC is connected. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SWITCH (warehouse name), Switch (caption), Switch (attribute name), and SWITCH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Transmitted attribute
The total transmission rate in KBps of this virtual machine NIC. The type is 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.

The following names are defined for this attribute: TRANSMITTED or T (warehouse name), Transmitted (caption), Transmitted (attribute name), and T (column name).

Usage attribute
The total rate in KBps that data is being transmitted and received data on this virtual machine NIC. The type is 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.

The following names are defined for this attribute: USAGE (warehouse name), Usage (caption), Usage (attribute name), and USAGE (column name).

Virtual Machine attribute
The name of the virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VIRTUAL_MACHINE or VM (warehouse name), Virtual Machine (caption), Virtual_Machine (attribute name), and VM (column name).

VM NIC attribute
The name of the virtual machine NIC. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_NIC (warehouse name), VM NIC (caption), VM_NIC (attribute name), and VM_NIC (column name).

Networked Virtual Switches attribute group

This attribute group contains information about the standard virtual switches in the virtual infrastructure grouped by network. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Datacenter attribute
The name of the data center that uses this virtual switch. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Number_Of_NICs < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Transmitted < 0) || (Received < 0) || (Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Managed System Name attribute
The managed system name of the subnode for the ESX server of the virtual switch. The type is 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.

The following names are defined for this attribute: SUBNODE_MSN or SM (warehouse name), Managed System Name (caption), Subnode_MSN (attribute name), and SM (column name).

Network attribute
The name of the network with which the virtual switch is associated. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NETWORK (warehouse name), Network (caption), Network (attribute name), and NETWORK (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Number NICs attribute
The number of NICs connected to the virtual switch. The type is 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.

The following names are defined for this attribute: NUMBER_OF_NICS or NON (warehouse name), Number NICs (caption), Number_Of_NICs (attribute name), and NON (column name).

Received attribute
The total reception rate in KBps of the host on this virtual switch. The type is 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.

The following names are defined for this attribute: RECEIVED (warehouse name), Received (caption), Received (attribute name), and RECEIVED (column name).

Server Hostname attribute
The hostname of the ESX server to which the virtual switch belongs. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Switch attribute
The name of the virtual switch. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SWITCH (warehouse name), Switch (caption), Switch (attribute name), and SWITCH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Transmitted attribute
The total transmission rate in KBps of the host on this virtual switch. The type is 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.

The following names are defined for this attribute: TRANSMITTED or T (warehouse name), Transmitted (caption), Transmitted (attribute name), and T (column name).

Usage attribute
The total rate in KBps that the host is transmitting and receiving data on this virtual switch. The type is 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.

The following names are defined for this attribute: USAGE (warehouse name), Usage (caption), Usage (attribute name), and USAGE (column name).

Networks attribute group

This attribute group contains information about the networks in the virtual infrastructure. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Datacenter attribute
The name of the data center that uses this network. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Distributed Switch attribute
The name of the distributed virtual switch for this network, if applicable. The type is 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.

The following names are defined for this attribute: DISTRIBUTED_SWITCH or DS (warehouse name), Distributed Switch (caption), Distributed_Switch (attribute name), and DS (column name).

Network attribute
The name of the network. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NETWORK (warehouse name), Network (caption), Network (attribute name), and NETWORK (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Number Hosts attribute
The number of hosts connected to the network. The type is 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.

The following names are defined for this attribute: NUMBER_OF_HOSTS or NOH (warehouse name), Number Hosts (caption), Number_Of_Hosts (attribute name), and NOH (column name).

Number VMs attribute
The number of virtual machines connected to the network. The type is 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.

The following names are defined for this attribute: NUMBER_OF_VMS or NOV (warehouse name), Number VMs (caption), Number_Of_VMs (attribute name), and NOV (column name).

Overall Status attribute
The overall alarm status of the network. A value of red or yellow indicates that an alarm has been triggered for the network. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), red (red), yellow (yellow), green (green), gray (gray). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Type attribute
The type of network. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NETWORK_TYPE or NT (warehouse name), Type (caption), Network_Type (attribute name), and NT (column name).

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

This attribute group contains the following attributes:

Average Collection Duration attribute
The average duration of all data collections of this group in seconds. The type is real number (32-bit counter) with two decimal places of precision with enumerated values. The following values are defined: NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AVERAGE_COLLECTION_DURATION or COLAVGD (warehouse name), Average Collection Duration (caption), Average_Collection_Duration (attribute name), and COLAVGD (column name).

Cache Hit Percent attribute
The percentage of external data requests for this group that were satisfied from the cache. The type is real number (32-bit counter) with two decimal places of precision.

The following names are defined for this attribute: CACHE_HIT_PERCENT or CACHPCT (warehouse name), Cache Hit Percent (caption), Cache_Hit_Percent (attribute name), and CACHPCT (column name).

Cache Hits attribute
The number of times an external data request for this group was satisfied from the cache. The type is integer (32-bit counter).

The following names are defined for this attribute: CACHE_HITS or CACHEHT (warehouse name), Cache Hits (caption), Cache_Hits (attribute name), and CACHEHT (column name).

Cache Misses attribute
The number of times an external data request for this group was not available in the cache. The type is integer (32-bit counter).

The following names are defined for this attribute: CACHE_MISSES or CACHEMS (warehouse name), Cache Misses (caption), Cache_Misses (attribute name), and CACHEMS (column name).

Error Code attribute

The error code associated with the query. The type is 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), KVM NO DATASOURCES (1000), KVM DATASOURCE LOGIN FAILED (1005), KVM DATASOURCE NOT FOUND (1010), SUBNODE UNAVAILABLE (1033), KVM PROVIDER RESET (2222). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ERROR_CODE or ERRCODE (warehouse name), Error Code (caption), Error_Code (attribute name), and ERRCODE (column name).

Intervals Skipped attribute

The number of times a background data collection for this group was skipped because the previous collection was still running when the next one was due to start. The type is integer (32-bit counter).

The following names are defined for this attribute: INTERVALS_SKIPPED or INTSKIP (warehouse name), Intervals Skipped (caption), Intervals_Skipped (attribute name), and INTSKIP (column name).

Last Collection Duration attribute

The duration of the most recently completed data collection of this group in seconds. The type is real number (32-bit counter) with two decimal places of precision.

The following names are defined for this attribute: LAST_COLLECTION_DURATION or COLDURA (warehouse name), Last Collection Duration (caption), Last_Collection_Duration (attribute name), and COLDURA (column name).

Last Collection Finished attribute

The most recent time a data collection of this group finished. The type is timestamp with enumerated values. The following values are defined: NOT COLLECTED (0691231190000000), NOT COLLECTED (0000000000000001). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_COLLECTION_FINISHED or COLFINI (warehouse name), Last Collection Finished (caption), Last_Collection_Finished (attribute name), and COLFINI (column name).

Last Collection Start attribute

The most recent time a data collection of this group started. The type is timestamp with enumerated values. The following values are defined: NOT COLLECTED (0691231190000000), NOT COLLECTED

(0000000000000001). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LAST_COLLECTION_START or COLSTRT (warehouse name), Last Collection Start (caption), Last_Collection_Start (attribute name), and COLSTRT (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Number of Collections attribute
The number of data collections for this group since the agent started. The type is integer (32-bit counter).

The following names are defined for this attribute: NUMBER_OF_COLLECTIONS or NUMCOLL (warehouse name), Number of Collections (caption), Number_of_Collections (attribute name), and NUMCOLL (column name).

Object Name attribute
The name of the performance object. The type is string.

The following names are defined for this attribute: OBJECT_NAME or OBJNAME (warehouse name), Object Name (caption), Object_Name (attribute name), and OBJNAME (column name).

Object Status attribute
The status of the performance object. The type is 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.

The following names are defined for this attribute: OBJECT_STATUS or OBJSTTS (warehouse name), Object Status (caption), Object_Status (attribute name), and OBJSTTS (column name).

Object Type attribute
The type of the performance object. The type is 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.

The following names are defined for this attribute: OBJECT_TYPE or OBJTYPE (warehouse name), Object Type (caption), Object_Type (attribute name), and OBJTYPE (column name).

Query Name attribute
The name of the attribute group. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: QUERY_NAME or ATTRGRP (warehouse name), Query Name (caption), Query_Name (attribute name), and ATTRGRP (column name).

Refresh Interval attribute
The interval at which this group is refreshed in seconds. The type is integer (32-bit counter).

The following names are defined for this attribute: REFRESH_INTERVAL or REFRINT (warehouse name), Refresh Interval (caption), Refresh_Interval (attribute name), and REFRINT (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Resource Pool CPU attribute group

This attribute group contains information about CPU metrics for resource pools. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

CPU Usage attribute
The CPU usage in MHz of all running child virtual machines including virtual machines in child resource pools. The type is 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.

The following names are defined for this attribute: CPU_USAGE (warehouse name), CPU Usage (caption), CPU_Usage (attribute name), and CPU_USAGE (column name).

Expandable attribute
Indicates if the CPU reservation is permitted to grow beyond the specified configuration value when the parent resource pool has sufficient unreserved CPU resource. The type is integer with enumerated values. The following values are defined: Yes (1), No (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: EXPANDABLE (warehouse name), Expandable (caption), Expandable (attribute name), and EXPANDABLE (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Percent_Overall_Usage < 0) || (Max_Usage < 0) || (CPU_Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Reservation_Used < 0) || (Percent_Reserved_VMs < 0) || (Reservation_Used_VM < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Include Data In Summarization 3 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Unreserved < 0) || (Unreserved_VM < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_3 or IDIS3 (warehouse name), Include Data In Summarization 3 (caption), Include_Data_In_Summarization_3 (attribute name), and IDIS3 (column name).

Limit attribute
The configured upper limit of CPU resources in MHz that this resource pool can get even if there are sufficient resources that would otherwise permit the limit to be higher. A value of -1 indicates that there is no limit. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-2), No limit (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LIMIT (warehouse name), Limit (caption), Limit (attribute name), and LIMIT (column name).

Max Usage attribute
The current upper bound on CPU usage in MHz.This limit is based on the limit that is configured for the resource pool and the limits that are configured for all parent resource pools. The type is 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.

The following names are defined for this attribute: MAX_USAGE (warehouse name), Max Usage (caption), Max_Usage (attribute name), and MAX_USAGE (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Parent Name attribute
The name of the parent of this resource pool. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: PARENT_NAME or PN (warehouse name), Parent Name (caption), Parent_Name (attribute name), and PN (column name).

Percent Overall Usage attribute
The percentage of CPU resources being used relative to the maximum amount currently available. The type is 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.

The following names are defined for this attribute: PERCENT_OVERALL_USAGE or POU (warehouse name), Percent Overall Usage (caption), Percent_Overall_Usage (attribute name), and POU (column name).

Percent Reserved VMs attribute
The percentage of CPU resources that are reserved for virtual machines. The type is 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.

The following names are defined for this attribute: PERCENT_RESERVED_VMS or PRV (warehouse name), Percent Reserved VMs (caption), Percent_Reserved_VMs (attribute name), and PRV (column name).

Pool Name attribute
The name of this resource pool. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: POOL_NAME (warehouse name), Pool Name (caption), Pool_Name (attribute name), and POOL_NAME (column name).

Reservation attribute
The amount of CPU resource in MHz that is guaranteed to be available to the resource pool. The type is 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.

The following names are defined for this attribute: RESERVATION or R (warehouse name), Reservation (caption), Reservation (attribute name), and R (column name).

Reservation Used attribute
The total amount of CPU resources in MHz that have been used to satisfy the reservation requirements of all descendants of this resource pool. The type is 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.

The following names are defined for this attribute: RESERVATION_USED or RU (warehouse name), Reservation Used (caption), Reservation_Used (attribute name), and RU (column name).

Reservation Used VM attribute
The total amount of CPU resources in MHz that have been used to satisfy the reservations of running virtual machines in this resource pool and its descendants. The type is 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.

The following names are defined for this attribute: RESERVATION_USED_VM or RUV (warehouse name), Reservation Used VM (caption), Reservation_Used_VM (attribute name), and RUV (column name).

Server Hostname attribute
The host name of the ESX server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Share Level attribute
The named level for the defined number of shares. This level corresponds to the Shares attribute. The type is 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.

The following names are defined for this attribute: SHARE_LEVEL or SL (warehouse name), Share Level (caption), Share_Level (attribute name), and SL (column name).

Shares attribute
The relative weighting of CPU allocations given to this resource pool in actual numeric form. This attribute is only applicable when the shares level is defined as custom. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-2), Not applicable (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SHARES (warehouse name), Shares (caption), Shares (attribute name), and SHARES (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Unreserved attribute
The total amount of CPU resources in MHz available to satisfy a reservation for child resource pool. The type is 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.

The following names are defined for this attribute: UNRESERVED (warehouse name), Unreserved (caption), Unreserved (attribute name), and UNRESERVED (column name).

Unreserved VM attribute
The total amount of CPU resources available in MHz to satisfy a reservation for a child virtual machine. The type is 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.

The following names are defined for this attribute: UNRESERVED_VM or UV (warehouse name), Unreserved VM (caption), Unreserved_VM (attribute name), and UV (column name).

Resource Pool General attribute group

This attribute group contains information about general metrics and the configuration of resource pools. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

CPU Usage attribute
The CPU usage in MHz of all running child virtual machines including virtual machines in child resource pools. The type is 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.

The following names are defined for this attribute: CPU_USAGE (warehouse name), CPU Usage (caption), CPU_Usage (attribute name), and CPU_USAGE (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Number_Child_Pools < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (CPU_Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Memory_Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Memory Usage attribute
The memory usage in MB of all running child virtual machines including virtual machines in child resource pools. The type is 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.

The following names are defined for this attribute: MEMORY_USAGE or MU (warehouse name), Memory Usage (caption), Memory_Usage (attribute name), and MU (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Number Child Pools attribute
The number of resource pools that are immediate children of this resource pool. The type is 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.

The following names are defined for this attribute: NUMBER_CHILD_POOLS or NCP (warehouse name), Number Child Pools (caption), Number_Child_Pools (attribute name), and NCP (column name).

Number VMs attribute
The number of virtual machines that are children of this resource pool including virtual machines in child resource pools. The type is 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.

The following names are defined for this attribute: NUMBER_VMS (warehouse name), Number VMs (caption), Number_VMs (attribute name), and NUMBER_VMS (column name).

Number VMs On attribute
The number of virtual machines that are children of this resource pool including virtual machines in child resource pools that are powered on. The type is 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.

The following names are defined for this attribute: NUMBER_VMS_ON or NVO (warehouse name), Number VMs On (caption), Number_VMs_On (attribute name), and NVO (column name).

Overall Status attribute
The overall status indication of this resource pool. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), red (red), yellow (yellow), green (green), gray (gray). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STATUS (warehouse name), Overall Status (caption), Status (attribute name), and STATUS (column name).

Parent Name attribute
The name of the parent of this resource pool. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: PARENT_NAME or PN (warehouse name), Parent Name (caption), Parent_Name (attribute name), and PN (column name).

Pool Name attribute
The name of this resource pool. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: POOL_NAME (warehouse name), Pool Name (caption), Pool_Name (attribute name), and POOL_NAME (column name).

Server Hostname attribute
The host name of the ESX server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Resource Pool Memory attribute group

This attribute group contains information about memory metrics for resource pools. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Expandable attribute
Indicates if the memory reservation is permitted to grow beyond the specified configuration value when the parent resource pool has sufficient unreserved CPU resource. The type is integer with enumerated values. The following values are defined: Yes (1), No (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: EXPANDABLE (warehouse name), Expandable (caption), Expandable (attribute name), and EXPANDABLE (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Memory_Usage < 0) || (Max_Usage < 0) || (Percent_Overall_Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Reservation_Used < 0) || (Reservation_Used_VM < 0) || (Percent_Reserved_VMs < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Include Data In Summarization 3 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Unreserved < 0) || (Unreserved_VM < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_3 or IDIS3 (warehouse name), Include Data In Summarization 3 (caption), Include_Data_In_Summarization_3 (attribute name), and IDIS3 (column name).

Limit attribute
The configured upper limit of memory resources in MB that this resource pool can get even if there are sufficient resources that would otherwise permit the limit to be higher. A value of -1 indicates that there is no limit. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-2), No limit (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: LIMIT (warehouse name), Limit (caption), Limit (attribute name), and LIMIT (column name).

Max Usage attribute
The current upper bound on memory usage in MB. This value is based on the limit configured for this resource pool and the limits configured for all parent resource pools. The type is 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.

The following names are defined for this attribute: MAX_USAGE (warehouse name), Max Usage (caption), Max_Usage (attribute name), and MAX_USAGE (column name).

Memory Usage attribute
The memory usage in MB of all running child virtual machines including virtual machines in child resource pools. The type is 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.

The following names are defined for this attribute: MEMORY_USAGE or MU (warehouse name), Memory Usage (caption), Memory_Usage (attribute name), and MU (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Parent Name attribute
The name of the parent of this resource pool. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: PARENT_NAME or PN (warehouse name), Parent Name (caption), Parent_Name (attribute name), and PN (column name).

Percent Overall Usage attribute
The percentage of memory resources being used relative to the maximum amount currently available. The type is 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.

The following names are defined for this attribute: PERCENT_OVERALL_USAGE or POU (warehouse name), Percent Overall Usage (caption), Percent_Overall_Usage (attribute name), and POU (column name).

Percent Reserved VMs attribute
The percentage of memory resources that are reserved for virtual machines. The type is 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.

The following names are defined for this attribute: PERCENT_RESERVED_VMS or PRV (warehouse name), Percent Reserved VMs (caption), Percent_Reserved_VMs (attribute name), and PRV (column name).

Pool Name attribute
The name of this resource pool. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: POOL_NAME (warehouse name), Pool Name (caption), Pool_Name (attribute name), and POOL_NAME (column name).

Reservation attribute
The amount of memory resource in MB that is guaranteed to be available to the resource pool. The type is 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.

The following names are defined for this attribute: RESERVATION or R (warehouse name), Reservation (caption), Reservation (attribute name), and R (column name).

Reservation Used attribute
The total amount of memory resources in MB that have been used to satisfy the reservation requirements of all descendants of this resource pool. The type is 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.

The following names are defined for this attribute: RESERVATION_USED or RU (warehouse name), Reservation Used (caption), Reservation_Used (attribute name), and RU (column name).

Reservation Used VM attribute
The total amount of memory resources in MB that have been used to satisfy the reservations of running virtual machines in this resource pool and its descendants. The type is 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.

The following names are defined for this attribute: RESERVATION_USED_VM or RUV (warehouse name), Reservation Used VM (caption), Reservation_Used_VM (attribute name), and RUV (column name).

Server Hostname attribute
The host name of the ESX server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Share Level attribute
The named level for the defined number of shares. This value corresponds to the Shares attribute. The type is 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.

The following names are defined for this attribute: SHARE_LEVEL or SL (warehouse name), Share Level (caption), Share_Level (attribute name), and SL (column name).

Shares attribute
The relative weighting of memory allocations given to this resource pool. This attribute is applicable only when the shares level is defined as custom. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-2), Not applicable (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SHARES (warehouse name), Shares (caption), Shares (attribute name), and SHARES (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Unreserved attribute
The total amount of memory resources in MB available to satisfy a reservation for a child resource pool. The type is 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.

The following names are defined for this attribute: UNRESERVED (warehouse name), Unreserved (caption), Unreserved (attribute name), and UNRESERVED (column name).

Unreserved VM attribute
The total amount of memory resources available in MB to satisfy a reservation for a child virtual machine. The type is 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.

The following names are defined for this attribute: UNRESERVED_VM or UV (warehouse name), Unreserved VM (caption), Unreserved_VM (attribute name), and UV (column name).

Server attribute group

This attribute group contains basic information about an ESX server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Average VM CPU Percent Ready attribute
The average of all CPU percent ready values for all the virtual machines on this ESX server. The type is real number (32-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.

The following names are defined for this attribute: AVG_VM_CPU_PERCENT_RDY or AVCPR (warehouse name), Average VM CPU Percent Ready (caption), Avg_VM_CPU_Percent_Rdy (attribute name), and AVCPR (column name).

BIOS Date attribute
The date of release for this system's BIOS. The type is 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.

The following names are defined for this attribute: BIOS_DATE (warehouse name), BIOS Date (caption), BIOS_Date (attribute name), and BIOS_DATE (column name).

Build number attribute
The VMware product build number for the installed level of ESX. The type is 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.

The following names are defined for this attribute: BUILD_NUMBER or BN (warehouse name), Build number (caption), Build_number (attribute name), and BN (column name).

capacity attribute
The total capacity of the vFlash resource in MB(Megabytes) connected to this server. The type is 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.

The following names are defined for this attribute: CAPACITY (warehouse name), capacity (caption), Capacity (attribute name), and CAPACITY (column name).

Capacity For Vm Cache attribute
The total capacity of vFlash resource in MB(Megabytes) that can be allocated for VM caches. The type is 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.

The following names are defined for this attribute: CAPACITYFORVMCACHE or C (warehouse name), Capacity For Vm Cache (caption), CapacityForVmCache (attribute name), and C (column name).

CloneFrom Snapshot Supported attribute
Indicates whether cloning a virtual machine from a snapshot point is allowed. The type is integer with enumerated values. The following values are defined: True (1), False (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CLONEFROMSNAPSHOT_SUPPORTED or CS0 (warehouse name), CloneFrom Snapshot Supported (caption), CloneFromSnapshot_Supported (attribute name), and CS0 (column name).

Cluster attribute
The name of the cluster that this server is a member of or unavailable if not a member of any cluster. The type is 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.

The following names are defined for this attribute: CLUSTER (warehouse name), Cluster (caption), Cluster (attribute name), and CLUSTER (column name).

Connection State attribute
The connection state of the server. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), connected (connected), disconnected (disconnected), notResponding (notResponding). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CONNECTION_STATE or CS (warehouse name), Connection State (caption), Connection_State (attribute name), and CS (column name).

CPU Packages attribute
The number of CPU packages for this host. The type is 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.

The following names are defined for this attribute: CPU_PACKAGES or CP (warehouse name), CPU Packages (caption), CPU_Packages (attribute name), and CP (column name).

Current EVC Mode attribute
The current Enhanced VMotion Compatibility (EVC) mode of the host system. The type is 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.

The following names are defined for this attribute: CURRENT_EVC_MODE or CEM (warehouse name), Current EVC Mode (caption), Current_EVC_Mode (attribute name), and CEM (column name).

Datacenter attribute
The name of the data center for this server. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Datacenter MORef attribute
The internal managed object reference name of the data center for this server. The type is 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.

The following names are defined for this attribute: DATACENTER_MOREF or DM (warehouse name), Datacenter MORef (caption), Datacenter_MORef (attribute name), and DM (column name).

Datastore Space attribute
The total capacity in GB of the data stores connected to this server. This is across all of the data stores that this server is configured to use. The type is 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.

The following names are defined for this attribute: DATASTORE_SPACE or DS (warehouse name), Datastore Space (caption), Datastore_Space (attribute name), and DS (column name).

Datastore Used attribute
The total amount of datastore storage in GB that is actually in use by this server. This is across all of the data stores that this server is configured to use. The type is 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.

The following names are defined for this attribute: USED_DATASTORE or UD (warehouse name), Datastore Used (caption), Used_Datastore (attribute name), and UD (column name).

Demand attribute
The average active CPU load (in MHz) for the last minute. The type is 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.

The following names are defined for this attribute: DEMAND (warehouse name), Demand (caption), Demand (attribute name), and DEMAND (column name).

Energy Usage attribute
The amount of energy (in joules) that is used since the host system was started. The type is 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.

The following names are defined for this attribute: ENERGY_USAGE or EU (warehouse name), Energy Usage (caption), Energy_Usage (attribute name), and EU (column name).

Fault Tolerance Supported attribute
Indicates whether the host supports fault tolerance. The type is integer with enumerated values. The following values are defined: True (1), False (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FT_SUPPORTED or FS (warehouse name), Fault Tolerance Supported (caption), FT_Supported (attribute name), and FS (column name).

Fully Qualified Name attribute
This host's fully qualified name. The type is 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.

The following names are defined for this attribute: FULLY_QUALIFIED_NAME or FQN (warehouse name), Fully Qualified Name (caption), Fully_Qualified_Name (attribute name), and FQN (column name).

HBA Count attribute
The number of Host Bus Adapters (HBAs). The type is 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.

The following names are defined for this attribute: NUMBER_HBA (warehouse name), HBA Count (caption), Number_HBA (attribute name), and NUMBER_HBA (column name).

HyperThreading Enabled attribute
Whether hyperthreading is enabled on this server. The type is integer with enumerated values. The following values are defined: Yes (1), No (0), Unavailable (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: HYPERTHREADING_ENABLED or HE (warehouse name), HyperThreading Enabled (caption), HyperThreading_Enabled (attribute name), and HE (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Number_VMs < 0) || (Number_VMs_On < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (NICs < 0) || (Physical_CPUs < 0) || (CPU_Packages < 0) || (Physical_Memory < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Overall_CPU_Util < 0) || (Avg_VM_CPU_Percent_Rdy < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Include Data In Summarization 3 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Used_Datastore < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_3 or IDIS3 (warehouse name), Include Data In Summarization 3 (caption), Include_Data_In_Summarization_3 (attribute name), and IDIS3 (column name).

Include Data In Summarization 4 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Latency < 0) || (Demand < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_4 or IDIS4 (warehouse name), Include Data In Summarization 4 (caption), Include_Data_In_Summarization_4 (attribute name), and IDIS4 (column name).

Include Data In Summarization 5 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Energy_Usage < 0) || (Power_Usage < 0) || (Power_Capacity < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_5 or IDIS5 (warehouse name), Include Data In Summarization 5 (caption), Include_Data_In_Summarization_5 (attribute name), and IDIS5 (column name).

Include Data In Summarization 6 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Overall_Memory_Util < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_6 or IDIS6 (warehouse name), Include Data In Summarization 6 (caption), Include_Data_In_Summarization_6 (attribute name), and IDIS6 (column name).

Include Data In Summarization 7 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_VM_Configured_Memory < 0) || (Total_VM_Provisioned_Space < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_7 or IDIS7 (warehouse name), Include Data In Summarization 7 (caption), Include_Data_In_Summarization_7 (attribute name), and IDIS7 (column name).

Include Data In Summarization 8 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Storage_Adapter_Max_Latency < 0) || (Storage_Path_Max_Latency < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_8 or IDIS8 (warehouse name), Include Data In Summarization 8 (caption), Include_Data_In_Summarization_8 (attribute name), and IDIS8 (column name).

Include Data In Summarization 9 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Used_CPU_MHz < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_9 or IDIS9 (warehouse name), Include Data In Summarization 9 (caption), Include_Data_In_Summarization_9 (attribute name), and IDIS9 (column name).

IP Address attribute

The IP address of the host system. The type is 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.

The following names are defined for this attribute: IP_ADDRESS (warehouse name), IP Address (caption), IP_Address (attribute name), and IP_ADDRESS (column name).

Latency attribute

The amount of time (in percentage) that the resource pool waits in the ready state and is not scheduled because of a CPU resource contention. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: LATENCY (warehouse name), Latency (caption), Latency (attribute name), and LATENCY (column name).

Maintenance Mode attribute

Whether this server is in maintenance mode. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-1), No (0), Yes (1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAINTENANCE_MODE or MM (warehouse name), Maintenance Mode (caption), Maintenance_Mode (attribute name), and MM (column name).

Max EVC Mode attribute

The maximum Enhanced VMotion Compatibility (EVC) mode of the host system. The type is 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.

The following names are defined for this attribute: MAX_EVC_MODE or MEM (warehouse name), Max EVC Mode (caption), Max_EVC_Mode (attribute name), and MEM (column name).

NICs attribute

The number of NIC interfaces on this server. The type is 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.

The following names are defined for this attribute: NICS (warehouse name), NICs (caption), NICs (attribute name), and NICS (column name).

Node attribute

The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Number VMs attribute
The number of virtual machines configured on this server. The type is 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.

The following names are defined for this attribute: NUMBER_VMS (warehouse name), Number VMs (caption), Number_VMs (attribute name), and NUMBER_VMS (column name).

Number VMs On attribute
The number of virtual machines configured on this server that are powered on. The type is 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.

The following names are defined for this attribute: NUMBER_VMS_ON or NVO (warehouse name), Number VMs On (caption), Number_VMs_On (attribute name), and NVO (column name).

Overall CPU Util attribute
The overall CPU usage of the server. The type is 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.

The following names are defined for this attribute: OVERALL_CPU_UTIL or OCU (warehouse name), Overall CPU Util (caption), Overall_CPU_Util (attribute name), and OCU (column name).

Overall Memory Util attribute
The overall memory usage of the server. The type is 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.

The following names are defined for this attribute: OVERALL_MEMORY_UTIL or OMU (warehouse name), Overall Memory Util (caption), Overall_Memory_Util (attribute name), and OMU (column name).

Overall Status attribute
An indicator of the overall status of the server. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), red (red), yellow (yellow), green (green), gray (gray). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Performance Error Pct attribute
The percentage of performance monitoring API calls against this host that failed during their last execution. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: PERFORMANCE_ERROR_PCT or PEP (warehouse name), Performance Error Pct (caption), Performance_Error_Pct (attribute name), and PEP (column name).

Performance Error Rate attribute
The error rate of performance monitoring API calls against this host over a configured interval. The type is 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.

The following names are defined for this attribute: PERFORMANCE_ERROR_RATE or PER (warehouse name), Performance Error Rate (caption), Performance_Error_Rate (attribute name), and PER (column name).

Physical CPUs attribute
The number of physical CPUs on this server. The type is 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.

The following names are defined for this attribute: PHYSICAL_CPUS or PC (warehouse name), Physical CPUs (caption), Physical_CPUs (attribute name), and PC (column name).

Physical Memory attribute
The amount of physical memory on this server in MB. The type is 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.

The following names are defined for this attribute: PHYSICAL_MEMORY or PM (warehouse name), Physical Memory (caption), Physical_Memory (attribute name), and PM (column name).

Power Capacity attribute
The maximum amount of power (in watts) that can be used. The type is 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.

The following names are defined for this attribute: POWER_CAPACITY or PC0 (warehouse name), Power Capacity (caption), Power_Capacity (attribute name), and PC0 (column name).

Power State attribute
The power status of the host system.The valid values are POWERED_OFF, POWERED_ON, STAND_BY, and UNKNOWN. The type is 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.

The following names are defined for this attribute: POWER_STATE or PS (warehouse name), Power State (caption), Power_State (attribute name), and PS (column name).

Power Usage attribute
The amount of power (in watts) that is currently used. The type is 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.

The following names are defined for this attribute: POWER_USAGE or PU (warehouse name), Power Usage (caption), Power_Usage (attribute name), and PU (column name).

Processor Family attribute
The processor family of this host's CPUs. The type is 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.

The following names are defined for this attribute: PROCESSOR_FAMILY or PF (warehouse name), Processor Family (caption), Processor_Family (attribute name), and PF (column name).

Product attribute
The VMware product string for the installed level of ESX. The type is 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.

The following names are defined for this attribute: PRODUCT (warehouse name), Product (caption), Product (attribute name), and PRODUCT (column name).

Serial Number attribute
The serial number of the hardware of the host system. The type is 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.

The following names are defined for this attribute: SERIAL_NUMBER or SN (warehouse name), Serial Number (caption), Serial_Number (attribute name), and SN (column name).

Server Hostname attribute
The host name of the ESX server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

SSH Status attribute
The SSH service status of the ESX Server.The value Up indicates that the service is running.The value Down indicates that the service has stopped, Unknown Host indicates that the server is unreachable, and Unavailable indicates that the service status is unknown. The type is integer with enumerated values. The following values are defined: Unavailable (-1), Down (0), Up (1), Unknown Host (2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SSH_STATUS (warehouse name), SSH Status (caption), SSH_Status (attribute name), and SSH_STATUS (column name).

Storage Adapter Max Latency attribute
The highest latency (in milliseconds) across all the storage adapters that are used by the host. The type is 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.

The following names are defined for this attribute: STORAGE_ADAPTER_MAX_LATENCY or SAML (warehouse name), Storage Adapter Max Latency (caption), Storage_Adapter_Max_Latency (attribute name), and SAML (column name).

Storage Path Max Latency attribute
The highest latency (in milliseconds) across all the storage paths that are used by the host. The type is 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.

The following names are defined for this attribute: STORAGE_PATH_MAX_LATENCY or SPML (warehouse name), Storage Path Max Latency (caption), Storage_Path_Max_Latency (attribute name), and SPML (column name).

System Model attribute
The system model of this host. The type is 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.

The following names are defined for this attribute: SYSTEM_MODEL or SM (warehouse name), System Model (caption), System_Model (attribute name), and SM (column name).

System Up Time attribute
The number of seconds since the server was started. The type is default 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.

The following names are defined for this attribute: SYSTEM_UP_TIME or SUT (warehouse name), System Up Time (caption), System_up_time (attribute name), and SUT (column name).

System Vendor attribute
The system vendor of this host. The type is 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.

The following names are defined for this attribute: SYSTEM_VENDOR or SV (warehouse name), System Vendor (caption), System_Vendor (attribute name), and SV (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total CPU MHz attribute
The total amount of the CPU of the server in MHz. The type is 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.

The following names are defined for this attribute: TOTAL_CPU_MHZ or TCM (warehouse name), Total CPU MHz (caption), Total_CPU_MHz (attribute name), and TCM (column name).

Total VM Configured Memory attribute
The total amount of memory in GB configured for all VMs on this server, The type is real number (32-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.

The following names are defined for this attribute: TOTAL_VM_CONFIGURED_MEMORY or TVCM (warehouse name), Total VM Configured Memory (caption), Total_VM_Configured_Memory (attribute name), and TVCM (column name).

Total VM Provisioned Space attribute
The total amount of space in GB, that has been provisioned for use by VMs on this server. The type is real number (32-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.

The following names are defined for this attribute: TOTAL_VM_PROVISIONED_SPACE or TVPS (warehouse name), Total VM Provisioned Space (caption), Total_VM_Provisioned_Space (attribute name), and TVPS (column name).

usage attribute
The total usage of vFlash resource in MB(Megabytes). The type is 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.

The following names are defined for this attribute: USAGE (warehouse name), usage (caption), Usage (attribute name), and USAGE (column name).

Used CPU MHz attribute
The amount of the CPU (in MHz) that is used by the server. The type is 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.

The following names are defined for this attribute: USED_CPU_MHZ or UCM (warehouse name), Used CPU MHz (caption), Used_CPU_MHz (attribute name), and UCM (column name).

UUID attribute
The UUID of the server. The type is 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.

The following names are defined for this attribute: UUID (warehouse name), UUID (caption), UUID (attribute name), and UUID (column name).

Version attribute
The VMware product version for the installed level of ESX. The type is 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.

The following names are defined for this attribute: VERSION (warehouse name), Version (caption), Version (attribute name), and VERSION (column name).

vMotion enabled attribute
A flag to indicate whether vMotion is configured on this server. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), Yes (Yes), No (No). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: VMOTION_ENABLED or VE (warehouse name), vMotion enabled (caption), vMotion_enabled (attribute name), and VE (column name).

Server CPU attribute group

This attribute group contains information about CPU usage for a server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Core Utilization attribute
The percentage of the CPU core that is currently utilized.A core is utilized if either a single or both the logical CPU cores are utilized when hyper-threading is enabled. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: CORE_UTILIZATION or CU0 (warehouse name), Core Utilization (caption), Core_Utilization (attribute name), and CU0 (column name).

CPU Number attribute
The number of this CPU. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: CPU_NUMBER (warehouse name), CPU Number (caption), CPU_Number (attribute name), and CPU_NUMBER (column name).

CPU Utilization attribute
The usage of this CPU. The type is 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.

The following names are defined for this attribute: CPU_UTILIZATION or CU (warehouse name), CPU Utilization (caption), CPU_Utilization (attribute name), and CU (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (CPU_Utilization < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Core_Utilization < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Server Hostname attribute
The host name of the ESX server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Server DataStore attribute group

This attribute group contains information about data stores for a server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Agent MSN attribute
This attribute is only for IBM-internal use. This attribute is a key attribute. The type is 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.

The source for this attribute is derived: getenv("KQZ_INSTANCE_NAME") +":"+getenv("CTIRA_HOSTNAME")+":VM"

The following names are defined for this attribute: AGENT_MSN (warehouse name), Agent MSN (caption), Agent_MSN (attribute name), and AGENT_MSN (column name).

Capacity attribute
The storage capacity in MB. This metric does not apply to floppy or CD drives. The type is 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.

The following names are defined for this attribute: CAPACITY (warehouse name), Capacity (caption), Capacity (attribute name), and CAPACITY (column name).

Datacenter attribute
The name of this datacenter. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Datastore MORef attribute
The internal managed object reference name of this datastore. The type is 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.

The following names are defined for this attribute: DATASTORE_MOREF or DM (warehouse name), Datastore MORef (caption), Datastore_MORef (attribute name), and DM (column name).

Free Space attribute
The amount of available storage in MB. The type is 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.

The following names are defined for this attribute: FREE_SPACE (warehouse name), Free Space (caption), Free_Space (attribute name), and FREE_SPACE (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Free_Space < 0) || (Used_Space < 0) || (Percent_Used < 0) || (Percent_Free < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Read_Latency < 0) || (Write_Latency < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Maximum File Size attribute
The maximum size in KB of a file that might be allocated. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-1), > 2048GB (-2). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAXIMUM_FILE_SIZE or MFS (warehouse name), Maximum File Size (caption), Maximum_File_Size (attribute name), and MFS (column name).

Name attribute
The name of the data store. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NAME (warehouse name), Name (caption), Name (attribute name), and NAME (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Overall Status attribute
The overall status for this alarm. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Percent Free attribute
The percentage of unused space in this data store. The type is 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.

The following names are defined for this attribute: PERCENT_FREE or PF (warehouse name), Percent Free (caption), Percent_Free (attribute name), and PF (column name).

Percent Used attribute
The percentage of used space in the data store. The type is 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.

The following names are defined for this attribute: PERCENT_USED or PU (warehouse name), Percent Used (caption), Percent_Used (attribute name), and PU (column name).

Read Latency attribute
The average amount of time (in milliseconds) taken for a read operation from the datastore. The type is 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.

The following names are defined for this attribute: READ_LATENCY or RL (warehouse name), Read Latency (caption), Read_Latency (attribute name), and RL (column name).

Server Hostname attribute
The host name of the ESX server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Type attribute
The file system type of this data store: NFS or VMFS. The type is 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.

The following names are defined for this attribute: TYPE (warehouse name), Type (caption), Type (attribute name), and TYPE (column name).

Used Space attribute
The amount of allocated storage in MB. The type is 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.

The following names are defined for this attribute: USED_SPACE (warehouse name), Used Space (caption), Used_Space (attribute name), and USED_SPACE (column name).

Write Latency attribute
The average amount of time (in milliseconds) taken for a write operation from the datastore. The type is 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.

The following names are defined for this attribute: WRITE_LATENCY or WL (warehouse name), Write Latency (caption), Write_Latency (attribute name), and WL (column name).

Server Disk attribute group

This attribute group contains information about disk usage for a server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Backing Datastore attribute
The name of the data store that backs this server disk if there is one. The type is string with enumerated values. The following values are defined: Not Applicable (Not Applicable), Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BACKING_DATASTORE or BD (warehouse name), Backing Datastore (caption), Backing_Datastore (attribute name), and BD (column name).

BUS Resets attribute
The number of bus resets in the performance interval. The type is 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.

The following names are defined for this attribute: BUS_RESETS (warehouse name), BUS Resets (caption), BUS_Resets (attribute name), and BUS_RESETS (column name).

Commands attribute
The number of disk commands issued during the performance interval. The type is 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.

The following names are defined for this attribute: COMMANDS (warehouse name), Commands (caption), Commands (attribute name), and COMMANDS (column name).

Commands Aborted attribute
The number of disk commands stopped during the performance interval. The type is 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.

The following names are defined for this attribute: COMMANDS_ABORTED or CA (warehouse name), Commands Aborted (caption), Commands_Aborted (attribute name), and CA (column name).

Device Latency attribute
The average amount of time in milliseconds to complete an operation by the physical device. The type is 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.

The following names are defined for this attribute: DEVICE_LATENCY or DL (warehouse name), Device Latency (caption), Device_Latency (attribute name), and DL (column name).

Device Read Latency attribute
The average amount of time in milliseconds that a read operation took by the physical device. The type is 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.

The following names are defined for this attribute: DEVICE_READ_LATENCY or DRL (warehouse name), Device Read Latency (caption), Device_Read_Latency (attribute name), and DRL (column name).

Device Total Latency attribute
The sum of the average amount of time in milliseconds to complete read and write operations by the physical device. The type is 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.

The following names are defined for this attribute: DEVICE_TOTAL_LATENCY or DTL (warehouse name), Device Total Latency (caption), Device_Total_Latency (attribute name), and DTL (column name).

Device Write Latency attribute
The average amount of time in milliseconds that a write operation took by the physical device. The type is 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.

The following names are defined for this attribute: DEVICE_WRITE_LATENCY or DWL (warehouse name), Device Write Latency (caption), Device_Write_Latency (attribute name), and DWL (column name).

Disk Name attribute
The name of a virtual disk on the server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DISK_NAME (warehouse name), Disk Name (caption), Disk_Name (attribute name), and DISK_NAME (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Write < 0) || (Number_Write < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Device_Latency < 0) || (Device_Total_Latency < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Kernel_Latency < 0) || (Kernel_Total_Latency < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Include Data In Summarization 3 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Queue_Latency < 0) || (Queue_Total_Latency < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_3 or IDIS3 (warehouse name), Include Data In Summarization 3 (caption), Include_Data_In_Summarization_3 (attribute name), and IDIS3 (column name).

Include Data In Summarization 4 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_Read_Latency < 0) || (Device_Read_Latency < 0) || (Kernel_Read_Latency < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_4 or IDIS4 (warehouse name), Include Data In Summarization 4 (caption), Include_Data_In_Summarization_4 (attribute name), and IDIS4 (column name).

Include Data In Summarization 5 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_Write_Latency < 0) || (Device_Write_Latency < 0) || (Kernel_Write_Latency < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_5 or IDIS5 (warehouse name), Include Data In Summarization 5 (caption), Include_Data_In_Summarization_5 (attribute name), and IDIS5 (column name).

Include Data In Summarization 6 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_Latency < 0) || (Queue_Read_Latency < 0) || (Queue_Write_Latency < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_6 or IDIS6 (warehouse name), Include Data In Summarization 6 (caption), Include_Data_In_Summarization_6 (attribute name), and IDIS6 (column name).

Include Data In Summarization 7 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Read < 0) || (Number_Read < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_7 or IDIS7 (warehouse name), Include Data In Summarization 7 (caption), Include_Data_In_Summarization_7 (attribute name), and IDIS7 (column name).

Kernel Latency attribute

The average amount of time in milliseconds to complete an operation by the VMware kernel. The type is 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.

The following names are defined for this attribute: KERNEL_LATENCY or KL (warehouse name), Kernel Latency (caption), Kernel_Latency (attribute name), and KL (column name).

Kernel Read Latency attribute

The average amount of time in milliseconds that a read operation took by the VMware kernel. The type is 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.

The following names are defined for this attribute: KERNEL_READ_LATENCY or KRL (warehouse name), Kernel Read Latency (caption), Kernel_Read_Latency (attribute name), and KRL (column name).

Kernel Total Latency attribute
The sum of the average amount of time in milliseconds to complete read and write an operations by the VMware kernel. The type is 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.

The following names are defined for this attribute: KERNEL_TOTAL_LATENCY or KTL (warehouse name), Kernel Total Latency (caption), Kernel_Total_Latency (attribute name), and KTL (column name).

Kernel Write Latency attribute
The average amount of time in milliseconds that a write operation took by the VMware kernel. The type is 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.

The following names are defined for this attribute: KERNEL_WRITE_LATENCY or KWL (warehouse name), Kernel Write Latency (caption), Kernel_Write_Latency (attribute name), and KWL (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Number Read attribute
The number of read operations on the disk in the performance interval. The type is 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.

The following names are defined for this attribute: NUMBER_READ or NR (warehouse name), Number Read (caption), Number_Read (attribute name), and NR (column name).

Number Write attribute
The number of write operations on the disk in the performance interval. The type is 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.

The following names are defined for this attribute: NUMBER_WRITE or NW (warehouse name), Number Write (caption), Number_Write (attribute name), and NW (column name).

Queue Latency attribute
The average amount of time in milliseconds spent in the queue for the VMware kernel per IO command. The type is 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.

The following names are defined for this attribute: QUEUE_LATENCY or QL (warehouse name), Queue Latency (caption), Queue_Latency (attribute name), and QL (column name).

Queue Read Latency attribute
The average amount of time in milliseconds that a read operation spent in the queue for the VMware kernel. The type is 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.

The following names are defined for this attribute: QUEUE_READ_LATENCY or QRL (warehouse name), Queue Read Latency (caption), Queue_Read_Latency (attribute name), and QRL (column name).

Queue Total Latency attribute
The sum of the average amount of time in milliseconds spent in the queue for reads and writes in the VMware kernel per IO command. The type is 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.

The following names are defined for this attribute: QUEUE_TOTAL_LATENCY or QTL (warehouse name), Queue Total Latency (caption), Queue_Total_Latency (attribute name), and QTL (column name).

Queue Write Latency attribute
The average amount of time in milliseconds that a write operation spent in the queue for the VMware kernel. The type is 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.

The following names are defined for this attribute: QUEUE_WRITE_LATENCY or QWL (warehouse name), Queue Write Latency (caption), Queue_Write_Latency (attribute name), and QWL (column name).

Read attribute
The amount of data read in the interval for this disk in KBps. The type is 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.

The following names are defined for this attribute: READ (warehouse name), Read (caption), Read (attribute name), and READ (column name).

Server Hostname attribute
The host name of the ESX server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Latency attribute
The average total amount of time spent on an IO operation for both the physical and kernel layers. The type is 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.

The following names are defined for this attribute: TOTAL_LATENCY or TL (warehouse name), Total Latency (caption), Total_Latency (attribute name), and TL (column name).

Total Read Latency attribute
The average total amount of time in milliseconds spent on a read operation for both the physical and kernel layers. The type is 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.

The following names are defined for this attribute: TOTAL_READ_LATENCY or TRL (warehouse name), Total Read Latency (caption), Total_Read_Latency (attribute name), and TRL (column name).

Total Write Latency attribute
The average total amount of time in milliseconds spent on a write operation for both the physical and kernel layers. The type is 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.

The following names are defined for this attribute: TOTAL_WRITE_LATENCY or TWL (warehouse name), Total Write Latency (caption), Total_Write_Latency (attribute name), and TWL (column name).

Write attribute
The amount of data written in the interval for this disk in KB per second. The type is 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.

The following names are defined for this attribute: WRITE (warehouse name), Write (caption), Write (attribute name), and WRITE (column name).

Server HBA attribute group

This attribute group contains information about the host bus adapters (HBA) of the server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Bus attribute
The bus number of this HBA. The type is 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.

The following names are defined for this attribute: BUS (warehouse name), Bus (caption), Bus (attribute name), and BUS (column name).

Current Link Speed attribute
The current operating link speed (in megabits per second) of the port.This attribute is available for the HostInternetScsiHba HBA type. The type is 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.

The following names are defined for this attribute: CURRENT_LINK_SPEED or CLS (warehouse name), Current Link Speed (caption), Current_Link_Speed (attribute name), and CLS (column name).

Device attribute
The device name of this HBA. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DEVICE (warehouse name), Device (caption), Device (attribute name), and DEVICE (column name).

Driver attribute
The driver being used for this HBA. The type is 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.

The following names are defined for this attribute: DRIVER (warehouse name), Driver (caption), Driver (attribute name), and DRIVER (column name).

HBA Type attribute
The type of Host Bus Adapter (HBA).The valid values are HostBlockHba, HostFibreChannelHba, HostInternetScsiHba, and HostParallelScsiHba. The type is 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.

The following names are defined for this attribute: HBA_TYPE (warehouse name), HBA Type (caption), HBA_Type (attribute name), and HBA_TYPE (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Read < 0) || (Write < 0) || (Read_Latency < 0) || (Write_Latency < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Storage_Adapter_Throughput_Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Max Link Speed attribute
The maximum supported link speed (in megabits per second) of the port.This attribute is available for the HostInternetScsiHba HBA type. The type is 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.

The following names are defined for this attribute: MAX_LINK_SPEED or MLS (warehouse name), Max Link Speed (caption), Max_Link_Speed (attribute name), and MLS (column name).

Model attribute
The model string for this HBA. The type is 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.

The following names are defined for this attribute: MODEL (warehouse name), Model (caption), Model (attribute name), and MODEL (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

PCI ID attribute
The PCI ID for this HBA. The type is 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.

The following names are defined for this attribute: PCI_ID (warehouse name), PCI ID (caption), PCI_ID (attribute name), and PCI_ID (column name).

Read attribute
The average amount of data that is read (in KB per second) by the storage adapter. The type is 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.

The following names are defined for this attribute: READ (warehouse name), Read (caption), Read (attribute name), and READ (column name).

Read Latency attribute
The average amount of time (in milliseconds) over a given sample interval that the storage adapter consumes for a read operation to complete.This average amount of time is the sum of kernel latency and device latency. The type is 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.

The following names are defined for this attribute: READ_LATENCY or RL (warehouse name), Read Latency (caption), Read_Latency (attribute name), and RL (column name).

Server Hostname attribute
The host name of the ESX server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Speed attribute
The current operating speed (in KB per second) of the adapter.This attribute is available for the HostFibreChannelHba HBA type. The type is 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.

The following names are defined for this attribute: SPEED (warehouse name), Speed (caption), Speed (attribute name), and SPEED (column name).

Status attribute
The operational status for this HBA. The type is 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.

The following names are defined for this attribute: STATUS (warehouse name), Status (caption), Status (attribute name), and STATUS (column name).

Storage Adapter Throughput Usage attribute
The I/O rate (in KB per second) of the storage adapter. The type is 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.

The following names are defined for this attribute: STORAGE_ADAPTER_THROUGHPUT_USAGE or SATU (warehouse name), Storage Adapter Throughput Usage (caption), Storage_Adapter_Throughput_Usage (attribute name), and SATU (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Write attribute
The average amount of data that is written (in KB per second) by the storage adapter. The type is 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.

The following names are defined for this attribute: WRITE (warehouse name), Write (caption), Write (attribute name), and WRITE (column name).

Write Latency attribute
The average amount of time (in milliseconds) over a given sample interval that the storage adapter consumes for a write operation to complete.This average amount of time is the sum of kernel latency and device latency. The type is 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.

The following names are defined for this attribute: WRITE_LATENCY or WL (warehouse name), Write Latency (caption), Write_Latency (attribute name), and WL (column name).

Server Health attribute group

This attribute group contains ESX server health information. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Sensor Name attribute
The name of the sensor. The type is 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.

The following names are defined for this attribute: SENSOR_NAME or SN (warehouse name), Sensor Name (caption), Sensor_Name (attribute name), and SN (column name).

Sensor Status attribute
The operational status of the sensor. The type is 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.

The following names are defined for this attribute: SENSOR_STATUS or SS (warehouse name), Sensor Status (caption), Sensor_Status (attribute name), and SS (column name).

Sensor Type attribute
The type of sensor. The type is 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.

The following names are defined for this attribute: SENSOR_TYPE or ST (warehouse name), Sensor Type (caption), Sensor_Type (attribute name), and ST (column name).

Sensor Units attribute
The units of Sensor_Value. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), Not applicable (Not applicable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SENSOR_UNITS or SU (warehouse name), Sensor Units (caption), Sensor_Units (attribute name), and SU (column name).

Sensor Value attribute
The value of the sensor. The type is real number (32-bit gauge) with three decimal places of precision with enumerated values. The following values are defined: Unavailable (-2147483648), Not applicable (-2147483647). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: SENSOR_VALUE or SV (warehouse name), Sensor Value (caption), Sensor_Value (attribute name), and SV (column name).

Server Hostname attribute
The host name of the ESX server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Server Memory attribute group

This attribute group contains information about memory usage for a server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Active Memory attribute
The amount of memory that is actively used in KB. The type is 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.

The following names are defined for this attribute: ACTIVE_MEMORY or AM (warehouse name), Active Memory (caption), Active_Memory (attribute name), and AM (column name).

Active Write attribute
The amount of memory (in KB) that is written to disk. The type is 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.

The following names are defined for this attribute: ACTIVE_WRITE or AW (warehouse name), Active Write (caption), Active_Write (attribute name), and AW (column name).

Balloon Used attribute
The amount of memory used by the virtual machine memory control system in KB. The type is 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.

The following names are defined for this attribute: BALLOON_USED or BU (warehouse name), Balloon Used (caption), Balloon_Used (attribute name), and BU (column name).

Free Memory attribute
The amount of physical memory that is currently free in MB. The type is 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.

The following names are defined for this attribute: FREE_MEMORY or FM (warehouse name), Free Memory (caption), Free_Memory (attribute name), and FM (column name).

Granted Max Memory attribute
The maximum amount of memory (in KB) that can be used by the virtual machine. The type is 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.

The following names are defined for this attribute: GRANTED_MAX_MEMORY or GMM (warehouse name), Granted Max Memory (caption), Granted_Max_Memory (attribute name), and GMM (column name).

Granted Memory attribute
The amount of memory available for use in KB. The type is 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.

The following names are defined for this attribute: GRANTED_MEMORY or GM (warehouse name), Granted Memory (caption), Granted_Memory (attribute name), and GM (column name).

Granted Min Memory attribute
The minimum amount of memory (in KB) that can be used by the virtual machine. The type is 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.

The following names are defined for this attribute: GRANTED_MIN_MEMORY or GMM0 (warehouse name), Granted Min Memory (caption), Granted_Min_Memory (attribute name), and GMM0 (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Memory_Utilization < 0) || (Active_Memory < 0) || (Granted_Memory < 0) || (Granted_Max_Memory < 0) || (Granted_Min_Memory < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Swap_In_Rate < 0) || (Swap_Out_Rate < 0) || (Swap_Total_Rate < 0) || (Swap_Used < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Active_Write < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Include Data In Summarization 3 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Memory_Usage < 0) || (Free_Memory < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_3 or IDIS3 (warehouse name), Include Data In Summarization 3 (caption), Include_Data_In_Summarization_3 (attribute name), and IDIS3 (column name).

Include Data In Summarization 4 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Swap_In_Rate_Host_Cache < 0) || (Swap_Out_Rate_Host_Cache < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_4 or IDIS4 (warehouse name), Include Data In Summarization 4 (caption), Include_Data_In_Summarization_4 (attribute name), and IDIS4 (column name).

Low Free Threshold attribute

The threshold of the free host physical memory (in KB).The ESX server starts recovering the memory from the virtual machines by using ballooning and swapping when the threshold is reached. The type is 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.

The following names are defined for this attribute: LOW_FREE_THRESHOLD or LFT (warehouse name), Low Free Threshold (caption), Low_Free_Threshold (attribute name), and LFT (column name).

Memory Usage attribute

The amount of physical memory in use in MB. The type is 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.

The following names are defined for this attribute: MEMORY_USAGE or MU (warehouse name), Memory Usage (caption), Memory_Usage (attribute name), and MU (column name).

Memory Utilization attribute

The physical memory usage as a percentage of used physical memory divided by physical memory installed. The type is 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.

The following names are defined for this attribute: MEMORY_UTILIZATION or MU0 (warehouse name), Memory Utilization (caption), Memory_Utilization (attribute name), and MU0 (column name).

Node attribute

The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute

This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Physical Memory attribute

The amount of physical memory in MB on this server. The type is 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.

The following names are defined for this attribute: PHYSICAL_MEMORY or PM (warehouse name), Physical Memory (caption), Physical_Memory (attribute name), and PM (column name).

Server Hostname attribute
The host name of the ESX server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Service Console attribute
The amount of memory reserved by the service console for the server in KB. The type is 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.

The following names are defined for this attribute: SERVICE_CONSOLE or SC (warehouse name), Service Console (caption), Service_Console (attribute name), and SC (column name).

Swap In Rate attribute
The rate at which memory is swapped in in kilobytes per second. The type is 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.

The following names are defined for this attribute: SWAP_IN_RATE or SIR (warehouse name), Swap In Rate (caption), Swap_In_Rate (attribute name), and SIR (column name).

Swap In Rate From Host Cache attribute
The rate (in KB per second) at which the memory is swapped from the host cache to the active memory. The type is 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.

The following names are defined for this attribute: SWAP_IN_RATE_HOST_CACHE or SIRHC (warehouse name), Swap In Rate From Host Cache (caption), Swap_In_Rate_Host_Cache (attribute name), and SIRHC (column name).

Swap Out Rate attribute
The rate at which memory is swapped out in kilobytes per second. The type is 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.

The following names are defined for this attribute: SWAP_OUT_RATE or SOR (warehouse name), Swap Out Rate (caption), Swap_Out_Rate (attribute name), and SOR (column name).

Swap Out Rate From Host Cache attribute
The rate (in KB per second) at which the memory is swapped from the active memory to the host cache. The type is 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.

The following names are defined for this attribute: SWAP_OUT_RATE_HOST_CACHE or SORHC (warehouse name), Swap Out Rate From Host Cache (caption), Swap_Out_Rate_Host_Cache (attribute name), and SORHC (column name).

Swap Total Rate attribute
The total rate at which memory is swapped in or out in kilobytes per second. The type is 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.

The following names are defined for this attribute: SWAP_TOTAL_RATE or STR (warehouse name), Swap Total Rate (caption), Swap_Total_Rate (attribute name), and STR (column name).

Swap Used attribute
The amount of memory used for swap space in KB. The type is 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.

The following names are defined for this attribute: SWAP_USED (warehouse name), Swap Used (caption), Swap_Used (attribute name), and SWAP_USED (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Server Network attribute group

This attribute group contains information about network usage for a server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Cluster attribute
The name of the cluster that this ESX server is a member of or unavailable if not a member of any cluster. The type is 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.

The following names are defined for this attribute: CLUSTER (warehouse name), Cluster (caption), Cluster (attribute name), and CLUSTER (column name).

Datacenter attribute
The name of the data center this ESX Server is a member of. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Duplex attribute
The current operating mode of the NIC. The type is 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.

The following names are defined for this attribute: DUPLEX (warehouse name), Duplex (caption), Duplex (attribute name), and DUPLEX (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Usage < 0) || (Transmitted < 0) || (Received < 0) || (Pkts_Received < 0) || (Pkts_Transmitted < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Link_Utilization < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Transmit_Pkts_Dropped < 0) || (Receive_Pkts_Dropped < 0) || (Pkts_Dropped < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Link Speed attribute
The current operating speed of the NIC in MB per second (mbps). The type is 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.

The following names are defined for this attribute: LINK_SPEED (warehouse name), Link Speed (caption), Link_Speed (attribute name), and LINK_SPEED (column name).

Link Utilization attribute
The percent usage of the NIC relative to the capacity of the link (including duplex). The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: LINK_UTILIZATION or LU (warehouse name), Link Utilization (caption), Link_Utilization (attribute name), and LU (column name).

NIC Name attribute
The name or label of this network interface. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NIC_NAME (warehouse name), NIC Name (caption), NIC_Name (attribute name), and NIC_NAME (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Physical Address attribute
The physical address of this NIC. The type is 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.

The following names are defined for this attribute: PHYSICAL_ADDR or PA (warehouse name), Physical Address (caption), Physical_Addr (attribute name), and PA (column name).

Pkts Dropped attribute
The number of packets dropped in the performance interval. The type is 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.

The following names are defined for this attribute: PKTS_DROPPED or PD (warehouse name), Pkts Dropped (caption), Pkts_Dropped (attribute name), and PD (column name).

Pkts Received attribute
The number of packets received in the performance interval. The type is 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.

The following names are defined for this attribute: PKTS_RECEIVED or PR (warehouse name), Pkts Received (caption), Pkts_Received (attribute name), and PR (column name).

Pkts Transmitted attribute
The number of packets transmitted in the performance interval. The type is 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.

The following names are defined for this attribute: PKTS_TRANSMITTED or PT (warehouse name), Pkts Transmitted (caption), Pkts_Transmitted (attribute name), and PT (column name).

Receive Pkts Dropped attribute
The number of receive packets dropped in the performance interval. The type is 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.

The following names are defined for this attribute: RECEIVE_PKTS_DROPPED or RPD (warehouse name), Receive Pkts Dropped (caption), Receive_Pkts_Dropped (attribute name), and RPD (column name).

Received attribute
The amount of data received in the performance interval in KB per second. The type is 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.

The following names are defined for this attribute: RECEIVED (warehouse name), Received (caption), Received (attribute name), and RECEIVED (column name).

Server Hostname attribute
The host name of the ESX server. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Status attribute
The current status, up or down, of the NIC. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), down (down), up (up). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STATUS (warehouse name), Status (caption), Status (attribute name), and STATUS (column name).

Switch attribute
The name of the virtual switch that the NIC is configured with. The type is 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.

The following names are defined for this attribute: VIRTUAL_SWITCH or VS (warehouse name), Switch (caption), Virtual_Switch (attribute name), and VS (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Transmit Pkts Dropped attribute
The number of transmit packets dropped in the performance interval. The type is 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.

The following names are defined for this attribute: TRANSMIT_PKTS_DROPPED or TPD (warehouse name), Transmit Pkts Dropped (caption), Transmit_Pkts_Dropped (attribute name), and TPD (column name).

Transmitted attribute
The amount of data transmitted in the performance interval in KB per second. The type is 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.

The following names are defined for this attribute: TRANSMITTED or T (warehouse name), Transmitted (caption), Transmitted (attribute name), and T (column name).

Usage attribute
The sum of data transmitted and received in the performance interval in KB per second. The type is 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.

The following names are defined for this attribute: USAGE (warehouse name), Usage (caption), Usage (attribute name), and USAGE (column name).

Server SAN attribute group

This attribute group contains information about the SAN devices for a server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Broken Paths attribute
The number of broken paths the host has to the device. The type is 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.

The following names are defined for this attribute: BROKEN_PATHS or BP (warehouse name), Broken Paths (caption), Broken_Paths (attribute name), and BP (column name).

Datastore attribute
The name of the associated data store for the disk. The type is string with enumerated values. The following values are defined: Not applicable (Not applicable), Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DATASTORE (warehouse name), Datastore (caption), Datastore (attribute name), and DATASTORE (column name).

Disabled Paths attribute
The number of disabled paths the host has to the device. The type is 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.

The following names are defined for this attribute: DISABLED_PATHS or DP (warehouse name), Disabled Paths (caption), Disabled_Paths (attribute name), and DP (column name).

Disk Name attribute
The name of the disk. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DISK_NAME (warehouse name), Disk Name (caption), Disk_Name (attribute name), and DISK_NAME (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Path Selection Policy attribute
The path selection policy the host uses to determine how to access the device. The type is 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.

The following names are defined for this attribute: PATH_SELECTION_POLICY or PSP (warehouse name), Path Selection Policy (caption), Path_Selection_Policy (attribute name), and PSP (column name).

Paths attribute
The number of paths the host has to the device. The type is 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.

The following names are defined for this attribute: PATHS (warehouse name), Paths (caption), Paths (attribute name), and PATHS (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Server Virtual Switches attribute group

This attribute group contains information about the virtual switches in the virtual infrastructure. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Datacenter attribute
The name of the data center that uses this virtual switch. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Transmitted < 0) || (Received < 0) || (Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Network attribute
The name of the network with which the virtual switch is associated. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NETWORK (warehouse name), Network (caption), Network (attribute name), and NETWORK (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Number NICs attribute
The number of NICs connected to the virtual switch. The type is 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.

The following names are defined for this attribute: NUMBER_OF_NICS or NON (warehouse name), Number NICs (caption), Number_Of_NICs (attribute name), and NON (column name).

Received attribute
The amount of data received in the performance interval in KB per second. The type is 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.

The following names are defined for this attribute: RECEIVED (warehouse name), Received (caption), Received (attribute name), and RECEIVED (column name).

Server Hostname attribute
The host name of the ESX server that the virtual switch belongs to. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Switch attribute
The name of the virtual switch. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SWITCH (warehouse name), Switch (caption), Switch (attribute name), and SWITCH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Transmitted attribute
The amount of data transmitted in the performance interval in KB per second. The type is 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.

The following names are defined for this attribute: TRANSMITTED or T (warehouse name), Transmitted (caption), Transmitted (attribute name), and T (column name).

Usage attribute
The total usage of the virtual switch. The type is 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.

The following names are defined for this attribute: USAGE (warehouse name), Usage (caption), Usage (attribute name), and USAGE (column name).

Server VM Datastore Utilization attribute group

This attribute group contains information about how each virtual machine is utilizing a data store. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Committed attribute
The amount of space in GB, on this data store, that this virtual machine is using. The type is real number (32-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.

The following names are defined for this attribute: COMMITTED (warehouse name), Committed (caption), Committed (attribute name), and COMMITTED (column name).

DataCenter attribute
The name of the data center that contains this data store. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), DataCenter (caption), DataCenter (attribute name), and DATACENTER (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Committed < 0) || (Uncommitted < 0) || (Provisioned < 0) || (Unshared < 0) || (Percent_Committed < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Name attribute
The name of the data store. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NAME (warehouse name), Name (caption), Name (attribute name), and NAME (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Percent Committed attribute
The percentage of space on this datastore that is committed as a percentage of the provisioned amount. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: PERCENT_COMMITTED or PC (warehouse name), Percent Committed (caption), Percent_Committed (attribute name), and PC (column name).

Provisioned attribute
The total reserved amount of space in GB, on this data store, that this virtual machine can use. The type is real number (32-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.

The following names are defined for this attribute: PROVISIONED or P (warehouse name), Provisioned (caption), Provisioned (attribute name), and P (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Uncommitted attribute
The reserved but unused amount of space in GB, on this data store, that this virtual machine can use in the future. The type is real number (32-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.

The following names are defined for this attribute: UNCOMMITTED or U (warehouse name), Uncommitted (caption), Uncommitted (attribute name), and U (column name).

Unshared attribute
The amount of space in GB, on this data store, occupied by this virtual machine that is not shared with any other virtual machines. The type is real number (32-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.

The following names are defined for this attribute: UNSHARED (warehouse name), Unshared (caption), Unshared (attribute name), and UNSHARED (column name).

Virtual Machine attribute
The name of the virtual machine on the data store. The type is 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.

The following names are defined for this attribute: VIRTUAL_MACHINE or VM (warehouse name), Virtual Machine (caption), Virtual_Machine (attribute name), and VM (column name).

VMNodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: VMNODEID (warehouse name), VMNodeID (caption), VMNodeID (attribute name), and VMNODEID (column name).

SubNode Events attribute group

This attribute group contains events for a server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Category attribute
The severity level associated with the event by VMware. The type is 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.

The following names are defined for this attribute: CATEGORY (warehouse name), Category (caption), Category (attribute name), and CATEGORY (column name).

Compute Resource attribute
The compute resource associated with this event. The type is 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.

The following names are defined for this attribute: COMPUTE_RESOURCE or CR (warehouse name), Compute Resource (caption), Compute_Resource (attribute name), and CR (column name).

Datacenter attribute
The data center associated with this event. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Entity Type attribute
The entity type of the event. The type is 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.

The following names are defined for this attribute: ENTITY_TYPE or ET (warehouse name), Entity Type (caption), Entity_Type (attribute name), and ET (column name).

ESX Server UUID attribute
The UUID of the ESX server associated with this event. The type is 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.

The following names are defined for this attribute: ESX_SERVER_UUID or ESU (warehouse name), ESX Server UUID (caption), ESX_Server_UUID (attribute name), and ESU (column name).

Event attribute
The event data string. The type is 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.

The following names are defined for this attribute: EVENT (warehouse name), Event (caption), Event (attribute name), and EVENT (column name).

Event Seq Number attribute
A sequence number for the event. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: EVENT_SEQ_NUMBER or ESN (warehouse name), Event Seq Number (caption), Event_Seq_Number (attribute name), and ESN (column name).

Event Text attribute
The full event data string. The type is 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.

The following names are defined for this attribute: EVENT_TEXT (warehouse name), Event Text (caption), Event_Text (attribute name), and EVENT_TEXT (column name).

Event Time attribute
The time that the event occurred. The type is 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.

The following names are defined for this attribute: EVENT_TIME (warehouse name), Event Time (caption), Event_Time (attribute name), and EVENT_TIME (column name).

Event Type attribute
The type of event given by VMware. The type is 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.

The following names are defined for this attribute: EVENT_TYPE (warehouse name), Event Type (caption), Event_Type (attribute name), and EVENT_TYPE (column name).

Event Type ID attribute
The type ID of the event given by VMware. This is unavailable unless the event is an extended event. The type is 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.

The following names are defined for this attribute: EVENT_TYPE_ID or ETI (warehouse name), Event Type ID (caption), Event_Type_ID (attribute name), and ETI (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Event_Seq_Number < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Server Hostname attribute
The host name of the ESX server that originated this event. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Timestamp attribute
The time the event was generated. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

UserId attribute
The user ID that caused the event. The type is string with enumerated values. The following values are defined: Not applicable (Not applicable), Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: USERID (warehouse name), UserId (caption), UserId (attribute name), and USERID (column name).

Virtual Machine attribute
The virtual machine associated with this event. The type is 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.

The following names are defined for this attribute: VIRTUAL_MACHINE or VM (warehouse name), Virtual Machine (caption), Virtual_Machine (attribute name), and VM (column name).

Virtual Machine UUID attribute
The UUID of the virtual machine associated with this event. The type is 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.

The following names are defined for this attribute: VIRTUAL_MACHINE_UUID or VMU (warehouse name), Virtual Machine UUID (caption), Virtual_Machine_UUID (attribute name), and VMU (column name).

Tasks attribute group

This attribute group provides information about the tasks that are completed on the vCenter server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Completed Time attribute
The date and time when the task was completed. The type is 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.

The following names are defined for this attribute: COMPLETED_TIME or CT (warehouse name), Completed Time (caption), Completed_Time (attribute name), and CT (column name).

Error Message attribute
The reason for the task failure. The type is 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.

The following names are defined for this attribute: ERROR_MESSAGE or EM (warehouse name), Error Message (caption), Error_Message (attribute name), and EM (column name).

Initiated By attribute
The type of the entity that created the task.The valid values are user name, another schedule task name, alarm name, and system. The type is 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.

The following names are defined for this attribute: INITIATED_BY or IB (warehouse name), Initiated By (caption), Initiated_By (attribute name), and IB (column name).

Name attribute
The name of the task. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NAME (warehouse name), Name (caption), Name (attribute name), and NAME (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Queue Time attribute
The date and time when the task was created. The type is 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.

The following names are defined for this attribute: QUEUE_TIME (warehouse name), Queue Time (caption), Queue_Time (attribute name), and QUEUE_TIME (column name).

Source Hostname attribute
The host name of the data source that created the task. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SOURCE_HOSTNAME or SH (warehouse name), Source Hostname (caption), Source_Hostname (attribute name), and SH (column name).

Start Time attribute
The date and time when the task started running. The type is 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.

The following names are defined for this attribute: START_TIME (warehouse name), Start Time (caption), Start_Time (attribute name), and START_TIME (column name).

Status attribute
The status of the task.The valid values are error and success. The type is 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.

The following names are defined for this attribute: STATUS (warehouse name), Status (caption), Status (attribute name), and STATUS (column name).

Target Entity attribute
The name of the target managed entity for the task. The type is 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.

The following names are defined for this attribute: TARGET_ENTITY or TE (warehouse name), Target Entity (caption), Target_Entity (attribute name), and TE (column name).

Target Entity Type attribute
The type of the target managed entity for the task. The type is 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.

The following names are defined for this attribute: TARGET_ENTITY_TYPE or TET (warehouse name), Target Entity Type (caption), Target_Entity_Type (attribute name), and TET (column name).

Timestamp attribute
The time the event was generated. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Thread Pool Status attribute group

The Thread Pool Status attribute group contains information that reflects the status of the internal thread pool used to collect data asynchronously. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Node attribute

The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Thread Pool Active Threads attribute

The number of threads in the thread pool currently active doing work. The type is integer (32-bit gauge) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: THREAD_POOL_ACTIVE_THREADS or TPACTTH (warehouse name), Thread Pool Active Threads (caption), Thread_Pool_Active_Threads (attribute name), and TPACTTH (column name).

Thread Pool Avg Active Threads attribute

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

The following names are defined for this attribute: THREAD_POOL_AVG_ACTIVE_THREADS or TPAVGAT (warehouse name), Thread Pool Avg Active Threads (caption), Thread_Pool_Avg_Active_Threads (attribute name), and TPAVGAT (column name).

Thread Pool Avg Job Wait attribute

The average time a job spends waiting on the thread pool queue in seconds. The type is real number (32-bit gauge) with two decimal places of precision with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: THREAD_POOL_AVG_JOB_WAIT or TPAVJBW (warehouse name), Thread Pool Avg Job Wait (caption), Thread_Pool_Avg_Job_Wait (attribute name), and TPAVJBW (column name).

Thread Pool Avg Queue Length attribute

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

The following names are defined for this attribute: THREAD_POOL_AVG_QUEUE_LENGTH or TPAVGQL (warehouse name), Thread Pool Avg Queue Length (caption), Thread_Pool_Avg_Queue_Length (attribute name), and TPAVGQL (column name).

Thread Pool Max Active Threads attribute

The peak number of threads in the thread pool that were simultaneously active doing work. The type is integer (32-bit counter) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: THREAD_POOL_MAX_ACTIVE_THREADS or TPMAXAT (warehouse name), Thread Pool Max Active Threads (caption), Thread_Pool_Max_Active_Threads (attribute name), and TPMAXAT (column name).

Thread Pool Max Queue Length attribute
The peak length the thread pool queue reached. The type is integer (32-bit counter) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: THREAD_POOL_MAX_QUEUE_LENGTH or TPMAXQL (warehouse name), Thread Pool Max Queue Length (caption), Thread_Pool_Max_Queue_Length (attribute name), and TPMAXQL (column name).

Thread Pool Max Size attribute
The maximum number of threads that are allowed to exist in the thread pool. The type is integer (32-bit numeric property) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: THREAD_POOL_MAX_SIZE or TPMAXSZ (warehouse name), Thread Pool Max Size (caption), Thread_Pool_Max_Size (attribute name), and TPMAXSZ (column name).

Thread Pool Min Active Threads attribute
The smallest number of threads in the thread pool that were simultaneously active doing work. The type is integer (32-bit counter) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: THREAD_POOL_MIN_ACTIVE_THREADS or TPMINAT (warehouse name), Thread Pool Min Active Threads (caption), Thread_Pool_Min_Active_Threads (attribute name), and TPMINAT (column name).

Thread Pool Min Queue Length attribute
The minimum length the thread pool queue reached. The type is integer (32-bit counter) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: THREAD_POOL_MIN_QUEUE_LENGTH or TPMINQL (warehouse name), Thread Pool Min Queue Length (caption), Thread_Pool_Min_Queue_Length (attribute name), and TPMINQL (column name).

Thread Pool Queue Length attribute
The number of jobs currently waiting in the thread pool queue. The type is integer (32-bit gauge) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: THREAD_POOL_QUEUE_LENGTH or TPQLGTH (warehouse name), Thread Pool Queue Length (caption), Thread_Pool_Queue_Length (attribute name), and TPQLGTH (column name).

Thread Pool Size attribute
The number of threads currently existing in the thread pool. The type is integer (32-bit gauge) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: THREAD_POOL_SIZE or THPSIZE (warehouse name), Thread Pool Size (caption), Thread_Pool_Size (attribute name), and THPSIZE (column name).

Thread Pool Total Jobs attribute
The number of jobs that are completed by all threads in the pool since agent start. The type is integer (32-bit counter) with enumerated values. The following values are defined: NO DATA (-1), NO DATA (-100). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: THREAD_POOL_TOTAL_JOBS or TPTJOBS (warehouse name), Thread Pool Total Jobs (caption), Thread_Pool_Total_Jobs (attribute name), and TPTJOBS (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Topological Events attribute group

This attribute group posts events when ESX servers and virtual machines are created or destroyed, or when virtual machines are relocated using vMotion. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

DATASTORE UUID attribute
The Universal Unique ID of the data store that is associated with the event. The type is 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.

The following names are defined for this attribute: DATASTORE_UUID or DU (warehouse name), DATASTORE UUID (caption), DATASTORE_UUID (attribute name), and DU (column name).

Entity Type attribute
The type of topological entity to which the event applies. The type is string with enumerated values. The following values are defined: Host System (Host System), Virtual Machine (Virtual Machine). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ENTITY_TYPE or ET (warehouse name), Entity Type (caption), Entity_Type (attribute name), and ET (column name).

Event Type attribute
The type of topological event that occurred. The type is string with enumerated values. The following values are defined: Created (Created), Destroyed (Destroyed), Relocated (Relocated). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: EVENT_TYPE (warehouse name), Event Type (caption), Event_Type (attribute name), and EVENT_TYPE (column name).

Host UUID attribute
The UUID of the host system associated with this event. The type is 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.

The following names are defined for this attribute: HOST_UUID (warehouse name), Host UUID (caption), Host_UUID (attribute name), and HOST_UUID (column name).

Managed System Name attribute
The managed system name associated with this event. The type is 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.

The following names are defined for this attribute: MSN (warehouse name), Managed System Name (caption), MSN (attribute name), and MSN (column name).

Name attribute
The name of the virtual machine or host that is producing this toplogy update. The type is 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.

The following names are defined for this attribute: NAME (warehouse name), Name (caption), Name (attribute name), and NAME (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Server Hostname attribute
The host name of the ESX server that is associated with the event. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Timestamp attribute
The time the event was generated. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

VM UUID attribute
The UUID of the virtual machine associated with this event. The type is 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.

The following names are defined for this attribute: VM_UUID (warehouse name), VM UUID (caption), VM_UUID (attribute name), and VM_UUID (column name).

Topology attribute group

This attribute group contains information about the topology of servers and virtual machines. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

ConnectionType attribute
The connection type from this node to the parent of this node. The type is string.

The following names are defined for this attribute: CONNECTIONTYPE or C0 (warehouse name), ConnectionType (caption), ConnectionType (attribute name), and C0 (column name).

ConnectToNode attribute
Indicates a connection from the NodeID to the node specified here. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: CONNECTTONODE or C (warehouse name), ConnectToNode (caption), ConnectToNode (attribute name), and C (column name).

Datacenter attribute
The name of this data center. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Managed System Name attribute
The managed system name that is associated with the data. The type is 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.

The following names are defined for this attribute: MSN (warehouse name), Managed System Name (caption), MSN (attribute name), and MSN (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
The identifier for this node in the topology. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

NodeName attribute
The name of this node in the tree. The type is string.

The following names are defined for this attribute: NODENAME (warehouse name), NodeName (caption), NodeName (attribute name), and NODENAME (column name).

NodeStatus attribute
The status of this node. The type is string.

The following names are defined for this attribute: NODESTATUS (warehouse name), NodeStatus (caption), NodeStatus (attribute name), and NODESTATUS (column name).

NodeType attribute
The type of node in the tree. The type is string.

The following names are defined for this attribute: NODETYPE (warehouse name), NodeType (caption), NodeType (attribute name), and NODETYPE (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Triggered Alarms attribute group

This attribute group contains information about the alarms in the virtual infrastructure. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Affected Entity attribute
The name of the entity that was affected by this alarm. The type is 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.

The following names are defined for this attribute: AFFECTED_ENTITY or EN (warehouse name), Affected Entity (caption), Affected_Entity (attribute name), and EN (column name).

Alarm Name attribute
The name of the alarm that got triggered. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: ALARM_NAME (warehouse name), Alarm Name (caption), Alarm_Name (attribute name), and ALARM_NAME (column name).

Alarm Status attribute
The alarm status for this alarm. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), red (red), yellow (yellow), green (green), gray (gray). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: ALARM_STATUS or OS (warehouse name), Alarm Status (caption), Alarm_Status (attribute name), and OS (column name).

Alarm Triggered Time attribute
The time that this alarm is triggered. The type is 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.

The following names are defined for this attribute: ALARM_TRIGGERED_TIME or ATT (warehouse name), Alarm Triggered Time (caption), Alarm_Triggered_Time (attribute name), and ATT (column name).

Datacenter attribute
The name of this data center. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Description attribute
The description of this alarm. The type is 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.

The following names are defined for this attribute: DESCRIPTION or D (warehouse name), Description (caption), Description (attribute name), and D (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Triggered Entity attribute
The name of the entity that this alarm was triggered on. The type is 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.

The following names are defined for this attribute: TRIGGERED_ENTITY or TE (warehouse name), Triggered Entity (caption), Triggered_Entity (attribute name), and TE (column name).

Triggered Entity Type attribute
The type of the entity for which this alarm was triggered. The type is 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.

The following names are defined for this attribute: TRIGGERED_ENTITY_TYPE or TET (warehouse name), Triggered Entity Type (caption), Triggered_Entity_Type (attribute name), and TET (column name).

vCenters attribute group

This attribute group displays basic information about VMware data sources. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Agent Connection attribute
The current connection status of this agent to the configured data source. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-1), Down (0), Up (1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: AGENT_CONNECTION or AC (warehouse name), Agent Connection (caption), Agent_Connection (attribute name), and AC (column name).

Average CU Execution Time attribute
The number of seconds that the previously executed collection units executed, divided by the number of previously executed collection units. The type is real number (32-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.

The following names are defined for this attribute: AVERAGE_CU_EXECUTION_TIME or ACET (warehouse name), Average CU Execution Time (caption), Average_CU_Execution_Time (attribute name), and ACET (column name).

Average CU Queue Time attribute
The number of seconds that the previously queued collection units were queued, divided by the number of previously queued collection units. The type is real number (32-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.

The following names are defined for this attribute: AVERAGE_CU_QUEUE_TIME or ACQT (warehouse name), Average CU Queue Time (caption), Average_CU_Queue_Time (attribute name), and ACQT (column name).

Collection Units attribute
The total number of collection units. The type is 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.

The following names are defined for this attribute: COLLECTION_UNITS or CU (warehouse name), Collection Units (caption), Collection_Units (attribute name), and CU (column name).

Configured Address attribute
The host address of the data source as entered in the agent data source configuration. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: CONFIGURED_ADDRESS or CA (warehouse name), Configured Address (caption), Configured_Address (attribute name), and CA (column name).

Current CU Execution Time attribute
The number of seconds that the currently executing collection units have been executing, divided by the number of currently executing collection units. The type is real number (32-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.

The following names are defined for this attribute: CURRENT_CU_EXECUTION_TIME or CCET (warehouse name), Current CU Execution Time (caption), Current_CU_Execution_Time (attribute name), and CCET (column name).

Current CU Queue Time attribute
The number of seconds that the currently queued collection units have been queued, divided by the number of currently queued collection units. The type is real number (32-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.

The following names are defined for this attribute: CURRENT_CU_QUEUE_TIME or CCQT (warehouse name), Current CU Queue Time (caption), Current_CU_Queue_Time (attribute name), and CCQT (column name).

Executing Collection Units attribute
The total number of collection units currently executing. The type is 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.

The following names are defined for this attribute: EXECUTING_COLLECTION_UNITS or ECU (warehouse name), Executing Collection Units (caption), Executing_Collection_Units (attribute name), and ECU (column name).

FQDN attribute
The fully qualified domain name of the data source. The type is 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.

The following names are defined for this attribute: FQDN (warehouse name), FQDN (caption), FQDN (attribute name), and FQDN (column name).

Inventory Age attribute
The number of seconds elapsed since the last time the inventory was updated. The type is real number (32-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.

The following names are defined for this attribute: INVENTORY_AGE or IA (warehouse name), Inventory Age (caption), Inventory_Age (attribute name), and IA (column name).

IP Address attribute
The IP address of the data source. The type is 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.

The following names are defined for this attribute: IP_ADDRESS (warehouse name), IP Address (caption), IP_Address (attribute name), and IP_ADDRESS (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Queued Collection Units attribute
The total number of collection units currently queued. The type is 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.

The following names are defined for this attribute: QUEUED_COLLECTION_UNITS or QCU (warehouse name), Queued Collection Units (caption), Queued_Collection_Units (attribute name), and QCU (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Type attribute
The type of data source, which can be vCenter or ESX server. The type is string with enumerated values. The following values are defined: Unavailable (-1), ESX (0), vCenter (1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TYPE (warehouse name), Type (caption), Type (attribute name), and TYPE (column name).

Web Services Port attribute
The port through which the agent communicates with the data source. The type is 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.

The following names are defined for this attribute: WEB_SERVICES_PORT or WSP (warehouse name), Web Services Port (caption), Web_Services_Port (attribute name), and WSP (column name).

Virtual Machines attribute group

This attribute group contains basic information about the virtual machines running on a server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Cluster attribute
The name of the cluster that this virtual machine is a member of or unavailable if not a member of any cluster. The type is 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.

The following names are defined for this attribute: CLUSTER (warehouse name), Cluster (caption), Cluster (attribute name), and CLUSTER (column name).

Connection State attribute
The connection status of the virtual machine.The valid values are connected, disconnected, inaccessible, invalid, and orphaned. The type is 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.

The following names are defined for this attribute: CONNECTION_STATE or CS (warehouse name), Connection State (caption), Connection_State (attribute name), and CS (column name).

Consolidation Needed attribute
Indicates whether any disk of the virtual machine requires consolidation. The type is integer with enumerated values. The following values are defined: Unavailable (-1), No (0), Yes (1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CONSOLIDATIONNEEDED or C (warehouse name), Consolidation Needed (caption), consolidationNeeded (attribute name), and C (column name).

CPU Limit attribute
The CPU limit of the virtual machine in mhz. The type is 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.

The following names are defined for this attribute: CPU_LIMIT (warehouse name), CPU Limit (caption), CPU_Limit (attribute name), and CPU_LIMIT (column name).

CPU Reservation attribute
Minimum amount of CPU in mhz guaranteed to be available to the virtual machine. The type is 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.

The following names are defined for this attribute: CPU_RESERVATION or CR (warehouse name), CPU Reservation (caption), CPU_Reservation (attribute name), and CR (column name).

CPU Shares attribute

The number of CPU shares, the relative weight, allocated to this virtual machine. This number is the actual value when the shares level has been configured as 'custom'. In general, the more shares a virtual machine has the more resource it gets. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-2), Not applicable (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CPU_SHARES (warehouse name), CPU Shares (caption), CPU_Shares (attribute name), and CPU_SHARES (column name).

CPU Utilization attribute

The overall CPU usage of this virtual machine during the collection interval. The type is 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.

The following names are defined for this attribute: CPU_UTILIZATION or CU (warehouse name), CPU Utilization (caption), CPU_Utilization (attribute name), and CU (column name).

Datacenter attribute

The name of this data center. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Datastore Percent Utilization attribute

The overall datastore usage of this virtual machine during the collection interval. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: VM_DATASTOREPERCENT_UTILIZATION or VDU (warehouse name), Datastore Percent Utilization (caption), VM_DatastorePercent_Utilization (attribute name), and VDU (column name).

Fault Tolerance attribute

An indication of the protection of the virtual machine against hardware failures. This attribute can be configured with a secondary virtual machine or it can be running on a server that is a member of a cluster that is configured for High Availability. The type is integer with enumerated values. The following values are defined: None (-1), FT Primary (1), FT NonPrimary (2), HA (3). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: FAULT_TOLERANCE or FT (warehouse name), Fault Tolerance (caption), Fault_Tolerance (attribute name), and FT (column name).

FQDN attribute

The fully qualified domain name of the virtual machine. The type is 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.

The following names are defined for this attribute: FQDN (warehouse name), FQDN (caption), FQDN (attribute name), and FQDN (column name).

FT Instance UUID attribute

The instance UUID of the fault tolerance virtual machine. The type is 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.

The following names are defined for this attribute: FT_INSTANCE_UUID or FIU (warehouse name), FT Instance UUID (caption), FT_Instance_UUID (attribute name), and FIU (column name).

Guest OS Managed System Name attribute

The managed system name of the guest OS agent within the virtual machine. The type is 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.

The following names are defined for this attribute: GUESTOS_MSN or GM (warehouse name), Guest OS Managed System Name (caption), GuestOS_MSN (attribute name), and GM (column name).

Guest State attribute

The operational state of the guest operating system installed in this virtual machine. The values can be running, shuttingdown, resetting, standby, notrunning, and unknown. The type is 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.

The following names are defined for this attribute: GUEST_STATE or GS (warehouse name), Guest State (caption), Guest_State (attribute name), and GS (column name).

GuestOS Name attribute

The full name of the guest operating system for this virtual machine. The type is 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.

The following names are defined for this attribute: GUESTOS_NAME or GN (warehouse name), GuestOS Name (caption), GuestOS_Name (attribute name), and GN (column name).

Heartbeats attribute

The number of heartbeats received from the virtual machine. The type is default 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.

The following names are defined for this attribute: HEARTBEATS (warehouse name), Heartbeats (caption), Heartbeats (attribute name), and HEARTBEATS (column name).

Hostname attribute

The host name of this virtual machine. The type is 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.

The following names are defined for this attribute: HOSTNAME (warehouse name), Hostname (caption), Hostname (attribute name), and HOSTNAME (column name).

Include Data In Summarization 0 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Num_CPUs < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (CPU_Utilization < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Include Data In Summarization 3 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (VM_Percent_RDY < 0) || (Used_CPU_MHz < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_3 or IDIS3 (warehouse name), Include Data In Summarization 3 (caption), Include_Data_In_Summarization_3 (attribute name), and IDIS3 (column name).

Include Data In Summarization 4 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Number_Of_Snapshots < 0) || (Snapshot_Storage_Consumed < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_4 or IDIS4 (warehouse name), Include Data In Summarization 4 (caption), Include_Data_In_Summarization_4 (attribute name), and IDIS4 (column name).

Include Data In Summarization 5 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Memory_Limit < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_5 or IDIS5 (warehouse name), Include Data In Summarization 5 (caption), Include_Data_In_Summarization_5 (attribute name), and IDIS5 (column name).

Include Data In Summarization 6 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (CPU_Limit < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_6 or IDIS6 (warehouse name), Include Data In Summarization 6 (caption), Include_Data_In_Summarization_6 (attribute name), and IDIS6 (column name).

Instance UUID attribute

The virtual center specific 128-bit Universal Unique ID (UUID) of a virtual machine.The UUID is represented as a hexadecimal string.This identifier is used by VirtualCenter to uniquely identify all the virtual machine instances. The type is 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.

The following names are defined for this attribute: INSTANCE_UUID or IU (warehouse name), Instance UUID (caption), Instance_UUID (attribute name), and IU (column name).

IP Address attribute
The IP address of this virtual machine. The type is 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.

The following names are defined for this attribute: IP_ADDRESS (warehouse name), IP Address (caption), IP_Address (attribute name), and IP_ADDRESS (column name).

Memory Limit attribute
The memory limit of the virtual machine in MB. The type is 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.

The following names are defined for this attribute: MEMORY_LIMIT or ML (warehouse name), Memory Limit (caption), Memory_Limit (attribute name), and ML (column name).

Memory Reservation attribute
Minimum amount of memory in MB guaranteed to be available to the virtual machine. The type is 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.

The following names are defined for this attribute: MEMORY_RESERVATION or MR (warehouse name), Memory Reservation (caption), Memory_Reservation (attribute name), and MR (column name).

Memory Shares attribute
The number of memory shares, the relative weight, allocated to this virtual machine. This number is the actual value when the shares level has been configured as 'custom'. In general, the more shares a virtual machine has the more resource it gets. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-2), Not applicable (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MEMORY_SHARES or MS0 (warehouse name), Memory Shares (caption), Memory_Shares (attribute name), and MS0 (column name).

Memory Size attribute
The memory size of the virtual machine in MB. The type is 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.

The following names are defined for this attribute: MEMORY_SIZE or MS (warehouse name), Memory Size (caption), Memory_Size (attribute name), and MS (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Num CPUs attribute
The number of CPUs configured for this virtual machine. The type is 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.

The following names are defined for this attribute: NUM_CPUS (warehouse name), Num CPUs (caption), Num_CPUs (attribute name), and NUM_CPUS (column name).

Number Disks attribute
The number of disks that are connected to the virtual machine. The type is integer (32-bit numeric property) 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.

The following names are defined for this attribute: NUMBER_OF_DISKS or NOD (warehouse name), Number Disks (caption), Number_Of_Disks (attribute name), and NOD (column name).

Number NICs attribute
The number of NICs that are connected to the virtual machine. The type is integer (32-bit numeric property) 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.

The following names are defined for this attribute: NUMBER_OF_NICS or NON (warehouse name), Number NICs (caption), Number_Of_NICs (attribute name), and NON (column name).

Number Of Snapshots attribute
The number of snapshots stored for this virtual machine. The type is 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.

The following names are defined for this attribute: NUMBER_OF_SNAPSHOTS or NOS (warehouse name), Number Of Snapshots (caption), Number_Of_Snapshots (attribute name), and NOS (column name).

Overall Status attribute
The overall status for this alarm. The type is 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.

The following names are defined for this attribute: OVERALL_STATUS or OS (warehouse name), Overall Status (caption), Overall_Status (attribute name), and OS (column name).

Power Status attribute
The current power status of the virtual machine. The type is 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.

The following names are defined for this attribute: POWER_STATUS or PS (warehouse name), Power Status (caption), Power_Status (attribute name), and PS (column name).

Resource Pool attribute
The name of the resource pool of which this virtual machine is a member. The type is 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.

The following names are defined for this attribute: RESOURCE_POOL or RP (warehouse name), Resource Pool (caption), Resource_Pool (attribute name), and RP (column name).

Snapshot Storage Consumed attribute
The amount of disk space (in MB) that is used by the virtual machine for the snapshots. The type is 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.

The following names are defined for this attribute: SNAPSHOT_STORAGE_CONSUMED or SSC (warehouse name), Snapshot Storage Consumed (caption), Snapshot_Storage_Consumed (attribute name), and SSC (column name).

Storage DRS Enable attribute
Indicates whether the Storage DRS is enabled. The type is integer with enumerated values. The following values are defined: True (1), False (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: STORAGE_DRS_ENABLE or SDE (warehouse name), Storage DRS Enable (caption), Storage_DRS_Enable (attribute name), and SDE (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Tools Status attribute
The operational status of the VMware VM Tools package in the guest operating system. The type is 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.

The following names are defined for this attribute: TOOLS_STATUS or TS (warehouse name), Tools Status (caption), Tools_Status (attribute name), and TS (column name).

Universally Unique Identifier attribute
The UUID (Universally Unique Identifier) for this virtual machine. The type is 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.

The following names are defined for this attribute: UUID (warehouse name), Universally Unique Identifier (caption), UUID (attribute name), and UUID (column name).

Up Time attribute
The number of seconds since the virtual machine was started. The type is default 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.

The following names are defined for this attribute: UP_TIME (warehouse name), Up Time (caption), Up_Time (attribute name), and UP_TIME (column name).

Used CPU MHz attribute
The total amount of CPU used by this virtual machine during the last sample period measured in MHz. The type is 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.

The following names are defined for this attribute: USED_CPU_MHZ or UCM (warehouse name), Used CPU MHz (caption), Used_CPU_MHz (attribute name), and UCM (column name).

Version attribute
The version string for this virtual machine. The type is 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.

The following names are defined for this attribute: VERSION (warehouse name), Version (caption), Version (attribute name), and VERSION (column name).

VM MORef attribute
The internal managed object reference name of the virtual machine. The type is 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.

The following names are defined for this attribute: MOREF (warehouse name), VM MORef (caption), MORef (attribute name), and MOREF (column name).

VM Name attribute
The user-defined display name of this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_NAME (warehouse name), VM Name (caption), VM_Name (attribute name), and VM_NAME (column name).

VM OS Type attribute
The guest family for the operating system. The type is 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.

The following names are defined for this attribute: VM_OS_TYPE (warehouse name), VM OS Type (caption), VM_OS_Type (attribute name), and VM_OS_TYPE (column name).

VM Percent Ready attribute
The CPU percent ready metric across all the virtual machine CPUs. The type is real number (32-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.

The following names are defined for this attribute: VM_PERCENT_RDY or VPR (warehouse name), VM Percent Ready (caption), VM_Percent_RDY (attribute name), and VPR (column name).

VM Server Name attribute
The host name of the ESX server that runs this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_SERVER_NAME or VSN (warehouse name), VM Server Name (caption), VM_Server_Name (attribute name), and VSN (column name).

VM Template attribute
Indicates whether this virtual machine is a template instead of a regular virtual machine. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-1), Yes (1), No (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: TEMPLATE (warehouse name), VM Template (caption), Template (attribute name), and TEMPLATE (column name).

Virtual Switches attribute group

This attribute group contains information about the standard virtual switches in the virtual infrastructure. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Datacenter attribute
The name of the data center that uses this virtual switch. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Number_Of_NICs < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Transmitted < 0) || (Received < 0) || (Usage < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Managed System Name attribute
The managed system name of the subnode for the ESX server of the virtual switch. The type is 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.

The following names are defined for this attribute: SUBNODE_MSN or SM (warehouse name), Managed System Name (caption), Subnode_MSN (attribute name), and SM (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Number NICs attribute
The number of NICs connected to the virtual switch. The type is 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.

The following names are defined for this attribute: NUMBER_OF_NICS or NON (warehouse name), Number NICs (caption), Number_Of_NICs (attribute name), and NON (column name).

Received attribute
The total reception rate in KBps of the host on this virtual switch. The type is 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.

The following names are defined for this attribute: RECEIVED (warehouse name), Received (caption), Received (attribute name), and RECEIVED (column name).

Server Hostname attribute
The hostname of the ESX server to which the virtual switch belongs. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SERVER_HOSTNAME or SH (warehouse name), Server Hostname (caption), Server_Hostname (attribute name), and SH (column name).

Switch attribute
The name of the virtual switch. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SWITCH (warehouse name), Switch (caption), Switch (attribute name), and SWITCH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Transmitted attribute
The total transmission rate in KBps of the host on this virtual switch. The type is 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.

The following names are defined for this attribute: TRANSMITTED or T (warehouse name), Transmitted (caption), Transmitted (attribute name), and T (column name).

Usage attribute
The total rate in KBps that the host is transmitting and receiving data on this virtual switch. The type is 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.

The following names are defined for this attribute: USAGE (warehouse name), Usage (caption), Usage (attribute name), and USAGE (column name).

VM CPU attribute group

This attribute group contains information about CPU usage for virtual machines that are powered on. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

CPU Number attribute
The virtual CPU number. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: CPU_NUMBER (warehouse name), CPU Number (caption), CPU_Number (attribute name), and CPU_NUMBER (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Wait_Time < 0) || (Ready_Time < 0) || (Used_Time < 0) || (Utilization < 0) || (Percent_Rdy < 0) || (User_Time < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Percent Ready attribute
The CPU ready time percentage. This value is calculated as the amount of time the VM spent in the ready state divided by the size of the sample interval. The type is 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.

The following names are defined for this attribute: PERCENT_RDY or PR (warehouse name), Percent Ready (caption), Percent_Rdy (attribute name), and PR (column name).

Ready Time attribute
The amount of time the CPU spent in the ready state in milliseconds. The type is 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.

The following names are defined for this attribute: READY_TIME (warehouse name), Ready Time (caption), Ready_Time (attribute name), and READY_TIME (column name).

Sys Time attribute
The amount of time the CPU spent in the system state in milliseconds. The type is 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.

The following names are defined for this attribute: SYS_TIME (warehouse name), Sys Time (caption), Sys_Time (attribute name), and SYS_TIME (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Used Time attribute
The amount of time the CPU used in milliseconds. The type is 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.

The following names are defined for this attribute: USED_TIME (warehouse name), Used Time (caption), Used_Time (attribute name), and USED_TIME (column name).

User Time attribute
The amount of time the CPU spent in the user (non_system) state in milliseconds. The type is 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.

The following names are defined for this attribute: USER_TIME (warehouse name), User Time (caption), User_Time (attribute name), and USER_TIME (column name).

Utilization attribute
The CPU usage percentage. This value is calculated as user time divided by the sum of used, ready, and wait times. The type is 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.

The following names are defined for this attribute: UTILIZATION or U (warehouse name), Utilization (caption), Utilization (attribute name), and U (column name).

VM HostName attribute
The host name of the virtual machine. The type is 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.

The following names are defined for this attribute: VM_HOSTNAME or VH (warehouse name), VM HostName (caption), VM_HostName (attribute name), and VH (column name).

VM Name attribute
The user-defined display name of this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_NAME (warehouse name), VM Name (caption), VM_Name (attribute name), and VM_NAME (column name).

VM Name CPU Number attribute
A concatenation of the VM Name and the CPU ID. The type is 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.

The following names are defined for this attribute: VM_NAME_CPU_NUMBER or VNCN (warehouse name), VM Name CPU Number (caption), VM_Name_CPU_Number (attribute name), and VNCN (column name).

VM OS Type attribute
The family for the guest operating system. The type is 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.

The following names are defined for this attribute: VM_OS_TYPE (warehouse name), VM OS Type (caption), VM_OS_Type (attribute name), and VM_OS_TYPE (column name).

VM Server Name attribute
The host name of the ESX server that runs this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_SERVER_NAME or VSN (warehouse name), VM Server Name (caption), VM_Server_Name (attribute name), and VSN (column name).

Wait Time attribute
The amount of time the CPU spent in the wait state in milliseconds. The type is 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.

The following names are defined for this attribute: WAIT_TIME (warehouse name), Wait Time (caption), Wait_Time (attribute name), and WAIT_TIME (column name).

VM Datastore Utilization attribute group

This attribute group contains information about the how each virtual machine is utilizing a data store. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Committed attribute
The amount of space in GB, on this datastore, that is being used by this virtual machine. The type is real number (32-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.

The following names are defined for this attribute: COMMITTED (warehouse name), Committed (caption), Committed (attribute name), and COMMITTED (column name).

DataCenter attribute
The name of the data center that contains this datastore. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), DataCenter (caption), DataCenter (attribute name), and DATACENTER (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Committed < 0) || (Uncommitted < 0) || (Provisioned < 0) || (Unshared < 0) || (Percent_Committed < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Total_Read_KBps < 0) || (Total_Write_KBps < 0) || (Total_IO_KBps < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Name attribute

The name of this datastore. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: NAME (warehouse name), Name (caption), Name (attribute name), and NAME (column name).

Node attribute

The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute

This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Percent Committed attribute

The percentage of space on this datastore that is committed as a percentage of the provisoned amount. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: PERCENT_COMMITTED or PC (warehouse name), Percent Committed (caption), Percent_Committed (attribute name), and PC (column name).

Provisioned attribute

The total reserved amount of space in GB, on this datastore, that can be used by this virtual machine. The type is real number (32-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.

The following names are defined for this attribute: PROVISIONED or P (warehouse name), Provisioned (caption), Provisioned (attribute name), and P (column name).

Timestamp attribute

The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total IO attribute
The sum of total kilobytes read and written per second by this vm from this datastore. The type is 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.

The following names are defined for this attribute: TOTAL_IO_KBPS or TIK (warehouse name), Total IO (caption), Total_IO_KBps (attribute name), and TIK (column name).

Total Read attribute
The total kilobytes read per second by this vm from this datastore. The type is 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.

The following names are defined for this attribute: TOTAL_READ_KBPS or TRK (warehouse name), Total Read (caption), Total_Read_KBps (attribute name), and TRK (column name).

Total Write attribute
The total kilobytes written per second by this vm from this datastore. The type is 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.

The following names are defined for this attribute: TOTAL_WRITE_KBPS or TWK (warehouse name), Total Write (caption), Total_Write_KBps (attribute name), and TWK (column name).

Uncommitted attribute
The reserved but unused amount of space in GB, on this datastore, that can be used in the future by this virtual machine. The type is real number (32-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.

The following names are defined for this attribute: UNCOMMITTED or U (warehouse name), Uncommitted (caption), Uncommitted (attribute name), and U (column name).

Unshared attribute
The amount of space in GB, on this datastore, occupied by this virtual machine that is not shared with any other virtual machines. The type is real number (32-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.

The following names are defined for this attribute: UNSHARED (warehouse name), Unshared (caption), Unshared (attribute name), and UNSHARED (column name).

UUID attribute
The UUID of the virtual machine associated with this datastore. This attribute is a key attribute. The type is string.

The following names are defined for this attribute: UUID (warehouse name), UUID (caption), UUID (attribute name), and UUID (column name).

Virtual Machine attribute
The name of the virtual machine on the datastore. The type is 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.

The following names are defined for this attribute: VIRTUAL_MACHINE or VM (warehouse name), Virtual Machine (caption), Virtual_Machine (attribute name), and VM (column name).

VM Disk attribute group

This attribute group contains information about disk usage for virtual machines. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Access attribute
The disk access (read or write). The type is 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.

The following names are defined for this attribute: ACCESS (warehouse name), Access (caption), Access (attribute name), and ACCESS (column name).

Backing data store attribute
The name of the data store that backs this disk. The type is string with enumerated values. The following values are defined: Not Applicable (Not Applicable), Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BACKING_DATASTORE or BD (warehouse name), Backing data store (caption), Backing_Datastore (attribute name), and BD (column name).

Capacity attribute
The capacity of the disk in MB. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-1), Not applicable (0). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CAPACITY (warehouse name), Capacity (caption), Capacity (attribute name), and CAPACITY (column name).

Connected attribute
Indicates whether the disk is currently connected to the virtual machine. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), Yes (Yes), No (No). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: CONNECTED (warehouse name), Connected (caption), Connected (attribute name), and CONNECTED (column name).

Description attribute
The disk label and description. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DESCRIPTION or D (warehouse name), Description (caption), Description (attribute name), and D (column name).

Disk Shares attribute
The number of disk shares, or the relative weight, allocated to this virtual machine. This is the actual value when the shares level has been configured as 'custom'. In general, the more shares a virtual machine has the more resource it gets. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-2), Not Applicable (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: DISK_SHARES or DS (warehouse name), Disk Shares (caption), Disk_Shares (attribute name), and DS (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Removable attribute
Indicates whether the disk is a removable disk. The type is string with enumerated values. The following values are defined: Unavailable (Unavailable), removable (removable), non-removable (non-removable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: REMOVABLE (warehouse name), Removable (caption), Removable (attribute name), and REMOVABLE (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

VM HostName attribute
The host name of the virtual machine. The type is 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.

The following names are defined for this attribute: VM_HOSTNAME or VH (warehouse name), VM HostName (caption), VM_HostName (attribute name), and VH (column name).

VM Name attribute
The user-defined display name of this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_NAME (warehouse name), VM Name (caption), VM_Name (attribute name), and VM_NAME (column name).

VM OS Type attribute
The guest family for the operating system. The type is 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.

The following names are defined for this attribute: VM_OS_TYPE (warehouse name), VM OS Type (caption), VM_OS_Type (attribute name), and VM_OS_TYPE (column name).

VM Server Name attribute
The host name of the ESX server that runs this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_SERVER_NAME or VSN (warehouse name), VM Server Name (caption), VM_Server_Name (attribute name), and VSN (column name).

VM Disk Performance attribute group

This attribute group provides information about the performance of the disks that are associated with the virtual machines. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Backing Datastore attribute
The name of the datastore that backs this VM disk if there is one. The type is string with enumerated values. The following values are defined: Not Applicable (Not Applicable), Unavailable (Unavailable). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: BACKING_DATASTORE or BD (warehouse name), Backing Datastore (caption), Backing_Datastore (attribute name), and BD (column name).

Disk Name attribute
The name of the disk. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DISK_NAME (warehouse name), Disk Name (caption), Disk_Name (attribute name), and DISK_NAME (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Read < 0) || (Write < 0) || (Number_Read < 0) || (Number_Write < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Number Read attribute
The number of times the data was read from the disk. The type is 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.

The following names are defined for this attribute: NUMBER_READ or NR (warehouse name), Number Read (caption), Number_Read (attribute name), and NR (column name).

Number Write attribute
The number of times the data was written to the disk. The type is 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.

The following names are defined for this attribute: NUMBER_WRITE or NW (warehouse name), Number Write (caption), Number_Write (attribute name), and NW (column name).

Read attribute
The amount of data that is read (in KB per second) from the disk during the collection interval. The type is 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.

The following names are defined for this attribute: READ (warehouse name), Read (caption), Read (attribute name), and READ (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Virtual Machine Name attribute
The name of the virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VIRTUAL_MACHINE or VM (warehouse name), Virtual Machine Name (caption), Virtual_Machine (attribute name), and VM (column name).

VM MORef attribute
The internal managed object reference name of the virtual machine. The type is 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.

The following names are defined for this attribute: MOREF (warehouse name), VM MORef (caption), MORef (attribute name), and MOREF (column name).

VMNodeID attribute
This attribute is only for IBM-internal use. Provides the interval node ID of the virtual machine. The type is string.

The following names are defined for this attribute: VMNODEID (warehouse name), VMNodeID (caption), VMNodeID (attribute name), and VMNODEID (column name).

Write attribute
The amount of data that is written (in KB per second) to the disk during the collection interval. The type is 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.

The following names are defined for this attribute: WRITE (warehouse name), Write (caption), Write (attribute name), and WRITE (column name).

VM Memory attribute group

This attribute group contains information about memory usage for virtual machines. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Active attribute
The amount of memory (in MB) that is actively used. The type is 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.

The following names are defined for this attribute: ACTIVE (warehouse name), Active (caption), Active (attribute name), and ACTIVE (column name).

Balloon Usage attribute
The amount of memory in KB being used by the VMware balloon driver. The type is 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.

The following names are defined for this attribute: BALLOON_USAGE or BU (warehouse name), Balloon Usage (caption), Balloon_Usage (attribute name), and BU (column name).

Datacenter attribute
The name of the data center that contains this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Granted attribute
The amount of memory (in MB) that is mapped to the virtual machine. The type is 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.

The following names are defined for this attribute: GRANTED (warehouse name), Granted (caption), Granted (attribute name), and GRANTED (column name).

Guest Free attribute

The amount of guest OS memory currently free in MB. This value is calculated as the difference between MemoryTotalSize and MemoryGuestUsage. The type is 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.

The following names are defined for this attribute: GUEST_FREE (warehouse name), Guest Free (caption), Guest_Free (attribute name), and GUEST_FREE (column name).

Guest Usage attribute

The amount of memory being used by the guest operating system in MB. The value can be between 0 and the configured memory size of the virtual machine. The type is 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.

The following names are defined for this attribute: GUEST_USAGE or GU (warehouse name), Guest Usage (caption), Guest_Usage (attribute name), and GU (column name).

Guest Util attribute

The percentage of memory (average) that was used by the guest running in this virtual machine over the past sample interval. This value is calculated as the percentage of MemoryGuestUsage over MemoryTotalSize. The type is 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.

The following names are defined for this attribute: GUEST_UTIL (warehouse name), Guest Util (caption), Guest_Util (attribute name), and GUEST_UTIL (column name).

Host Free attribute

The amount of virtual machine memory currently free in MB. This value is calculated as the difference between MemoryTotalSize and MemoryHostUsage. The type is 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.

The following names are defined for this attribute: HOST_FREE (warehouse name), Host Free (caption), Host_Free (attribute name), and HOST_FREE (column name).

Host Usage attribute

The amount of host (server) memory in MB that is currently being used by the virtual machine. The type is 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.

The following names are defined for this attribute: HOST_USAGE (warehouse name), Host Usage (caption), Host_Usage (attribute name), and HOST_USAGE (column name).

Host Util attribute

The percentage of memory (average) that was used by the virtual machine over the past sample interval. This value is calculated as the percentage of MemoryHostUsage over MemoryTotalSize. The type is 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.

The following names are defined for this attribute: HOST_UTIL (warehouse name), Host Util (caption), Host_Util (attribute name), and HOST_UTIL (column name).

Include Data In Summarization 0 attribute

This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Host_Usage < 0) || (Host_Util < 0) || (Guest_Usage < 0) || (Guest_Util < 0) || (Guest_Free < 0) || (Host_Free < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Include Data In Summarization 1 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Min_Alloc < 0) || (Usage < 0) || (Swap_To_File < 0) || (Balloon_Usage < 0) || (Active < 0) || (Shared < 0) || (Granted < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_1 or IDIS1 (warehouse name), Include Data In Summarization 1 (caption), Include_Data_In_Summarization_1 (attribute name), and IDIS1 (column name).

Include Data In Summarization 2 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_2 or IDIS2 (warehouse name), Include Data In Summarization 2 (caption), Include_Data_In_Summarization_2 (attribute name), and IDIS2 (column name).

Max Alloc attribute
Maximum amount of memory in MB that can be used by the virtual machine. The value is -1 if there is no limit. The type is integer (32-bit gauge) with enumerated values. The following values are defined: Unavailable (-2), No limit (-1). Any value that does not have a definition here is displayed in the User Interface.

The following names are defined for this attribute: MAX_ALLOC (warehouse name), Max Alloc (caption), Max_Alloc (attribute name), and MAX_ALLOC (column name).

Min Alloc attribute
Minimum amount of memory in MB guaranteed to be allocated to the virtual machine. The type is 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.

The following names are defined for this attribute: MIN_ALLOC (warehouse name), Min Alloc (caption), Min_Alloc (attribute name), and MIN_ALLOC (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Shared attribute
The amount of memory (in MB) that is shared with other virtual machines. The type is 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.

The following names are defined for this attribute: SHARED (warehouse name), Shared (caption), Shared (attribute name), and SHARED (column name).

Swap In Rate attribute
The rate (in KB per second) at which the memory is swapped from the disk to the active memory, over the past sample interval. The type is 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.

The following names are defined for this attribute: SWAP_IN_RATE or SIR (warehouse name), Swap In Rate (caption), Swap_In_Rate (attribute name), and SIR (column name).

Swap Out Rate attribute
The rate (in KB per second) at which memory is swapped from the active memory to the disk, over the past sample interval. The type is 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.

The following names are defined for this attribute: SWAP_OUT_RATE or SOR (warehouse name), Swap Out Rate (caption), Swap_Out_Rate (attribute name), and SOR (column name).

Swap To File attribute
The total amount of virtual machine memory that has been swapped out to the swap file in KB. The type is 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.

The following names are defined for this attribute: SWAP_TO_FILE or STF (warehouse name), Swap To File (caption), Swap_To_File (attribute name), and STF (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Total Size attribute
Total amount of memory allocated to the virtual machine in MB. The type is 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.

The following names are defined for this attribute: TOTAL_SIZE (warehouse name), Total Size (caption), Total_Size (attribute name), and TOTAL_SIZE (column name).

Usage attribute
The amount of memory (in percentage) that is used from the total configured or available memory. The type is real number (32-bit gauge) with two 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.

The following names are defined for this attribute: USAGE (warehouse name), Usage (caption), Usage (attribute name), and USAGE (column name).

VM HostName attribute
The host name of the virtual machine. The type is 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.

The following names are defined for this attribute: VM_HOSTNAME or VH (warehouse name), VM HostName (caption), VM_HostName (attribute name), and VH (column name).

VM Name attribute
The user-defined display name of this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_NAME (warehouse name), VM Name (caption), VM_Name (attribute name), and VM_NAME (column name).

VM OS Type attribute
The guest family for the operating system. The type is 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.

The following names are defined for this attribute: VM_OS_TYPE (warehouse name), VM OS Type (caption), VM_OS_Type (attribute name), and VM_OS_TYPE (column name).

VM Server Name attribute
The host name of the ESX server that runs this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_SERVER_NAME or VSN (warehouse name), VM Server Name (caption), VM_Server_Name (attribute name), and VSN (column name).

VM Network attribute group

This attribute group contains information about the network usage for the virtual machines on this ESX server. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Cluster attribute
The name of the cluster that this virtual machine is a member of or unavailable if not a member of any cluster. The type is 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.

The following names are defined for this attribute: CLUSTER (warehouse name), Cluster (caption), Cluster (attribute name), and CLUSTER (column name).

Datacenter attribute
The name of the data center this virtual machine is a member of. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), Datacenter (caption), Datacenter (attribute name), and DATACENTER (column name).

Description attribute
The description of this NIC. The type is 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.

The following names are defined for this attribute: DESCRIPTION or D (warehouse name), Description (caption), Description (attribute name), and D (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Transmitted < 0) || (Received < 0) || (Pkts_Trans < 0) || (Pkts_Recd < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Network attribute
The network name that the virtual NIC is associated with. The type is 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.

The following names are defined for this attribute: NETWORK_NAME or NN (warehouse name), Network (caption), Network_Name (attribute name), and NN (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Physical Address attribute
The physical address of this NIC. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: PHYSICAL_ADDR or PA (warehouse name), Physical Address (caption), Physical_Addr (attribute name), and PA (column name).

Pkts Received attribute
The number of packets received in the sample interval. The type is 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.

The following names are defined for this attribute: PKTS_RECD (warehouse name), Pkts Received (caption), Pkts_Recd (attribute name), and PKTS_RECD (column name).

Pkts Transmitted attribute
The number of packets transmitted in the sample interval. The type is 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.

The following names are defined for this attribute: PKTS_TRANS (warehouse name), Pkts Transmitted (caption), Pkts_Trans (attribute name), and PKTS_TRANS (column name).

Received attribute
The amount of data received in the sample interval in KB per second. The type is 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.

The following names are defined for this attribute: RECEIVED (warehouse name), Received (caption), Received (attribute name), and RECEIVED (column name).

Switch attribute
The name of the virtual switch that interface uses. The type is 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.

The following names are defined for this attribute: SWITCH (warehouse name), Switch (caption), Switch (attribute name), and SWITCH (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Transmitted attribute
The amount of data transmitted in the sample interval in KB per second. The type is 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.

The following names are defined for this attribute: TRANSMITTED or T (warehouse name), Transmitted (caption), Transmitted (attribute name), and T (column name).

VM HostName attribute
The host name of the virtual machine. The type is 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.

The following names are defined for this attribute: VM_HOSTNAME or VH (warehouse name), VM HostName (caption), VM_HostName (attribute name), and VH (column name).

VM Name attribute
The user-defined display name of this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_NAME (warehouse name), VM Name (caption), VM_Name (attribute name), and VM_NAME (column name).

VM OS Type attribute
The guest family for the operating system. The type is 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.

The following names are defined for this attribute: VM_OS_TYPE (warehouse name), VM OS Type (caption), VM_OS_Type (attribute name), and VM_OS_TYPE (column name).

VM Server Name attribute
The host name of the ESX server that runs this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_SERVER_NAME or VSN (warehouse name), VM Server Name (caption), VM_Server_Name (attribute name), and VSN (column name).

VM Orphaned Disk attribute group

This attribute group provides information about the orphaned virtual machine disks. It shows the orphaned disks of data store that are dedicated to vCenter; Data may not be valid if data store is shared across or managed by more than one vCenters. Data may not be valid if virtual machine is renamed and name mismatch occurred to the vmdk files created on the data store. Administrator is expected to ensure the data validity before taking any actions against the disks showed in this attribute group. IBM will not be responsible for any data loss. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

DataCenter attribute
The name of the data center that the data store belongs to. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATACENTER (warehouse name), DataCenter (caption), DataCenter (attribute name), and DATACENTER (column name).

Datastore attribute
The name of the data store that the orphaned virtual machine disk belongs to. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATASTORE (warehouse name), Datastore (caption), Datastore (attribute name), and DATASTORE (column name).

Datastore Cluster attribute
The name of the data store cluster that the data store belongs to. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DATASTORE_CLUSTER or DC (warehouse name), Datastore Cluster (caption), Datastore_Cluster (attribute name), and DC (column name).

File Path attribute
The path of the orphaned virtual machine disk. The type is 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.

The following names are defined for this attribute: FILE_PATH (warehouse name), File Path (caption), File_Path (attribute name), and FILE_PATH (column name).

File Size attribute
The size (in MB) of the orphaned virtual machine disk. The type is 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.

The following names are defined for this attribute: FILE_SIZE (warehouse name), File Size (caption), File_Size (attribute name), and FILE_SIZE (column name).

Last Modified attribute
The time when the orphaned virtual machine disk was last modified. The type is 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.

The following names are defined for this attribute: LAST_MODIFIED or LM (warehouse name), Last Modified (caption), Last_Modified (attribute name), and LM (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Owner attribute
The name of the owner of the orphaned virtual machine disk. The type is 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.

The following names are defined for this attribute: OWNER (warehouse name), Owner (caption), Owner (attribute name), and OWNER (column name).

Source attribute
The host name of the data source. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: SOURCE (warehouse name), Source (caption), Source (attribute name), and SOURCE (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

VM Partition attribute group

This attribute group contains information about disk partitions for virtual machines. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Capacity attribute
The size of the partition in MB. The type is 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.

The following names are defined for this attribute: CAPACITY (warehouse name), Capacity (caption), Capacity (attribute name), and CAPACITY (column name).

Description attribute
The description or label of this disk partition. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: DESCRIPTION or D (warehouse name), Description (caption), Description (attribute name), and D (column name).

Free Space attribute
The amount of unused space in MB. The type is 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.

The following names are defined for this attribute: FREE_SPACE (warehouse name), Free Space (caption), Free_Space (attribute name), and FREE_SPACE (column name).

Include Data In Summarization 0 attribute
This attribute is only for IBM-internal use.Indicates whether to include certain attribute data (numbers) in Tivoli Data Warehouse summarization.The valid values are 0 (exclude) and 1 (include). The type is integer (32-bit gauge).

The source for this attribute is derived: (Free_Space < 0) || (Used_Space < 0) || (Percent_Used < 0) || (Percent_Free < 0)? 0 : 1

The following names are defined for this attribute: INCLUDE_DATA_IN_SUMMARIZATION_0 or IDIS0 (warehouse name), Include Data In Summarization 0 (caption), Include_Data_In_Summarization_0 (attribute name), and IDIS0 (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

NodeID attribute
This attribute is only for IBM-internal use. The type is string.

The following names are defined for this attribute: NODEID (warehouse name), NodeID (caption), NodeID (attribute name), and NODEID (column name).

Percent Free attribute
The percentage of space on the partition is unallocated. The type is 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.

The following names are defined for this attribute: PERCENT_FREE or PF (warehouse name), Percent Free (caption), Percent_Free (attribute name), and PF (column name).

Percent Used attribute
The percentage usage of used space. The type is 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.

The following names are defined for this attribute: PERCENT_USED or PU (warehouse name), Percent Used (caption), Percent_Used (attribute name), and PU (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

Used Space attribute
The amount of space used in MB. The type is 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.

The following names are defined for this attribute: USED_SPACE (warehouse name), Used Space (caption), Used_Space (attribute name), and USED_SPACE (column name).

VM HostName attribute
The host name of the virtual machine. The type is 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.

The following names are defined for this attribute: VM_HOSTNAME or VH (warehouse name), VM HostName (caption), VM_HostName (attribute name), and VH (column name).

VM Name attribute
The user-defined display name of this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_NAME (warehouse name), VM Name (caption), VM_Name (attribute name), and VM_NAME (column name).

VM OS Type attribute
The guest family for the operating system. The type is 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.

The following names are defined for this attribute: VM_OS_TYPE (warehouse name), VM OS Type (caption), VM_OS_Type (attribute name), and VM_OS_TYPE (column name).

VM Server Name attribute
The host name of the ESX server that runs this virtual machine. This attribute is a key attribute. The type is 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.

The following names are defined for this attribute: VM_SERVER_NAME or VSN (warehouse name), VM Server Name (caption), VM_Server_Name (attribute name), and VSN (column name).

VM Snapshot attribute group

This attribute group is for IBM-internal use only. This attribute group is not eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

VM SnapshotFileLayout attribute group

This attribute group is for IBM-internal use only. This attribute group is not eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

VM Snapshots attribute group

This attribute group provides information about the snapshots for the virtual machines. This attribute group is eligible for use with Tivoli Data Warehouse.

This attribute group contains the following attributes:

Creation Time(Deprecated) attribute
This attribute is deprecated, alternative to use is Creation_Timestamp The type is 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.

The source for this attribute is derived: Creation_Date_I

The following names are defined for this attribute: CREATION_TIME or CT (warehouse name), Creation Time(Deprecated) (caption), Creation_Time (attribute name), and CT (column name).

Creation Timestamp attribute
The date and time when the snapshot was created.Timestamp format is MM/DD/YY HH:MM:SS The type is timestamp.

The source for this attribute is derived: Creation_Time_I

The following names are defined for this attribute: CREATION_TIMESTAMP or CT0 (warehouse name), Creation Timestamp (caption), Creation_Timestamp (attribute name), and CT0 (column name).

Description attribute
The description of the snapshot. The type is 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.

The source for this attribute is derived: Description_I

The following names are defined for this attribute: DESCRIPTION or D (warehouse name), Description (caption), Description (attribute name), and D (column name).

Node attribute
The managed system name of the agent. This attribute is a key attribute. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: NODE (warehouse name), Node (caption), ORIGINNODE (attribute name), and ORIGINNODE (column name).

Snapshot Age attribute
The age of the snapshot, in days. The type is 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.

The source for this attribute is derived: SnapshotAge_I

The following names are defined for this attribute: SNAPSHOT_AGE or SA (warehouse name), Snapshot Age (caption), Snapshot_Age (attribute name), and SA (column name).

Snapshot MORef attribute
The internal managed object reference name of the snapshot. This attribute is a key attribute. The type is 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.

The source for this attribute is derived: ManRef_I

The following names are defined for this attribute: SNAPSHOT_MOREF or SM (warehouse name), Snapshot MORef (caption), Snapshot_MORef (attribute name), and SM (column name).

Snapshot Name attribute
The name of the snapshot. The type is 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.

The source for this attribute is derived: Snapshot_Name_I

The following names are defined for this attribute: SNAPSHOT_NAME or SN (warehouse name), Snapshot Name (caption), Snapshot_Name (attribute name), and SN (column name).

Snapshot State attribute
Indicates whether snapshot of a virtual machine has taken by including its memory component. PoweredOff represents snapshot has taken in offline state with no memory inclusion, PoweredOn indicates snapshot has included memory component, Unavailable indicates state is unknown. This attribute is supported since vCenter/ESXi 6.0 onward. The type is string with enumerated values. The following values are defined: Unavailable (-1), PoweredOff (PoweredOff), PoweredOn (PoweredOn). Any value that does not have a definition here is displayed in the User Interface.

The source for this attribute is derived: Memory_Key_I

The following names are defined for this attribute: SNAPSHOT_STATE or SS (warehouse name), Snapshot State (caption), Snapshot_State (attribute name), and SS (column name).

Space Consumed attribute
The amount of disk space (in MB) that is used by the snapshot. The type is 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.

The source for this attribute is derived: Space_Consumed_I

The following names are defined for this attribute: SPACE_CONSUMED or SC (warehouse name), Space Consumed (caption), Space_Consumed (attribute name), and SC (column name).

Timestamp attribute
The local time at the agent when the data was collected. The type is string.

The source for this attribute is the agent

The following names are defined for this attribute: TIMESTAMP (warehouse name), Timestamp (caption), Timestamp (attribute name), and TIMESTAMP (column name).

VM Name attribute
The name of the virtual machine. The type is 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.

The source for this attribute is derived: VM_Name_I

The following names are defined for this attribute: VM_NAME (warehouse name), VM Name (caption), VM_Name (attribute name), and VM_NAME (column name).

VM State attribute
The state of the virtual machine when the snapshot was created. The type is 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.

The source for this attribute is derived: Virtual_Machine_State_I

The following names are defined for this attribute: VM_STATE (warehouse name), VM State (caption), VM_State (attribute name), and VM_STATE (column name).

Disk capacity planning for historical data

Disk capacity planning for a monitoring agent is a prediction of the amount of disk space to be consumed by the historical data in each attribute group that is collecting historical data. Required disk storage is an important factor when you are defining data collection rules and your strategy for historical data collection.

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

Table
Table name as it is displayed in the warehouse database, if the attribute group is configured to be written to the warehouse. The table name listed here corresponds to the table name in "Attribute groups for the monitoring agent" on page 15.

Attribute group
Name of the attribute group that is used to create the table in the warehouse database if it is short enough to fit in the table naming constraints of the database that is being used for the warehouse. The attribute group name listed here corresponds to the Warehouse table name in "Attribute groups for the monitoring agent" on page 15.

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

Database bytes per row (warehouse)
Estimate of the record length for detailed records that are written to the warehouse database, if the attribute group is configured to be written to the warehouse. Detailed records are records that have been uploaded from the agent for long-term historical data collection. This estimate can be used for warehouse disk-space planning purposes.

Aggregate bytes per row (warehouse)
Estimate of the record length for aggregate records that are written to the warehouse database, if the attribute group is configured to be written to the warehouse. Aggregate records are created by the

Summarization agent for attribute groups that have been configured for summarization. This estimate can be used for warehouse disk-space planning purposes.

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

Table 1. Capacity planning for historical data logged by the VMware VI agent

Table	Attribute group	Bytes per row (agent)	Database bytes per row (warehouse)	Aggregate bytes per row (warehouse)

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

Chapter 3. Situations

A situation is a logical expression involving one or more system conditions. Situations are used to monitor the condition of systems in your network. You can manage situations from the Tivoli Enterprise Portal by using the Situation Editor 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 VMware VI. 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 VMware VI 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” on page 171.

Predefined situations

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

Agent level Navigator items

- VMware VI
  - Not applicable
- Clusters
  - KVM_Cluster_Bad_Status
  - KVM_Cluster_CPU_Util_High

  - KVM_Cluster_Effective_CPU_Low
  - KVM_Cluster_Effective_Mem_Low
  - KVM_Cluster_Effective_Svrs_Low
  - KVM_Cluster_Memory_Util_High
- Datastores
  - KVM_Datastore_Bad_Status
  - KVM_Datastore_Inaccessible
  - KVM_Datastore_Usage_High
- Events
  - KVM_Cluster_Critical_Event
  - KVM_Datastore_Critical_Event
  - KVM_VM_Critical_Event
- Monitored Servers
  - KVM_Collection_Error
  - KVM_Collection_Time_Excessive
  - KVM_Connection_Failure
  - KVM_Host_System_Created
  - KVM_Host_System_Created2
  - KVM_Host_System_Destroyed
  - KVM_Host_System_Destroyed2
  - KVM_Inventory_Out_Of_Date
  - KVM_Take_Action_Failure
  - KVM_Virtual_Machine_Created
  - KVM_Virtual_Machine_Created2
  - KVM_Virtual_Machine_Destroyed
  - KVM_Virtual_Machine_Destroyed2
  - KVM_Virtual_Machine_Relocated
  - KVM_Virtual_Machine_Relocated2
- Networks
  - Not applicable

VMware VI (ESX) subnode

- VMware VI
  - Not applicable
- CPU
  - KVM_VM_CPU_Ready_High
  - KVM_VM_CPU_Util_High
- Disk
  - KVM_Server_Disk_Reads_High
  - KVM_Server_Disk_Writes_High
  - KVM_VM_Disk_Free_Low
- ESX Server
  - KVM_ESX_Server_Disconnected

  - KVM_Host_Server_Bad_Status
  - KVM_Server_CPU_Util_High
  - KVM_Server_Critical_Event
  - KVM_Server_Datastore_Free_Low
  - KVM_Server_HBA_Fault
  - KVM_Server_Memory_Util_High
  - KVM_Server_VMotion_Event
  - KVM_Server_VM_Critical_Event
- Memory
  - KVM_VM_Guest_Memory_Util_High
  - KVM_VM_Host_Memory_Util_High
- Network
  - KVM_Server_NIC_Down
  - KVM_Server_Receive_Rate_High
  - KVM_Server_Transmit_Rate_High
  - KVM_VM_Receive_Rate_High
  - KVM_VM_Transmit_Rate_High
- Resource Pools
  - KVM_Resource_Pool_CPU_High
  - KVM_Resource_Pool_Memory_High
- Virtual Machines
  - KVM_Snapshots_High
  - KVM_VM_Bad_Status
  - KVM_VM_Powered_Off

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.

VMware VI navigator item

No predefined situations are included for this navigator item.

Clusters navigator item

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

KVM_Cluster_Bad_Status situation

Description
The status of the cluster is not green.

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

Formula
```
*IF *VALUE KVM_CLUSTERS.Overall_Status *NE 'Unavailable' *AND *VALUE
KVM_CLUSTERS.Overall_Status *NE 'green'
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_Cluster_CPU_Util_High situation

Description
The CPU utilization of the cluster is high.

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

Formula
```
*IF *VALUE KVM_CLUSTERS.CPU_Utilization *GT 90
```

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
1 minute 30 seconds

Situation persistence
The number of times the condition specified by 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.

KVM_Cluster_Effective_CPU_Low situation

Description
The effective CPU amount of the cluster is low.

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

Formula

```
*IF *VALUE KVM_CLUSTERS.Percent_Effective_CPU *GE 0 *AND *VALUE
KVM_CLUSTERS.Percent_Effective_CPU *LT 50
```

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
1 minute 30 seconds

Situation persistence
The number of times the condition specified by 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.

KVM_Cluster_Effective_Mem_Low situation

Description
The effective memory of the cluster amount is low.

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

Formula

```
*IF *VALUE KVM_CLUSTERS.Percent_Effective_Memory *GE 0 *AND *VALUE
KVM_CLUSTERS.Percent_Effective_Memory *LT 50
```

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
1 minute 30 seconds

Situation persistence
The number of times the condition specified by 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.

KVM_Cluster_Effective_Svrs_Low situation

Description
The number of effective servers in the cluster is low.

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

Formula
```
*IF *VALUE KVM_CLUSTERS.Percent_Effective_Servers *GE 0 *AND *VALUE
KVM_CLUSTERS.Percent_Effective_Servers *LT 30
```

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
1 minute 30 seconds

Situation persistence
The number of times the condition specified by 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.

KVM_Cluster_Memory_Util_High situation

Description
The memory utilization of the cluster is high.

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

Formula
```
*IF *VALUE KVM_CLUSTERS.Memory_Utilization *GT 90
```

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
1 minute 30 seconds

Situation persistence
The number of times the condition specified by 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.

Datastores navigator item

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

KVM_Datastore_Bad_Status situation

Description
The status of the data store is not green.

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

Formula
```
*IF *VALUE KVM_DATASTORES.Overall_Status *NE 'Unavailable' *AND *VALUE
KVM_DATASTORES.Overall_Status *NE 'green'
```

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 condition specified by 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.

KVM_Datastore_Inaccessible situation

Description
The connectivity status of the data store is currently false.

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

Formula
```
*IF *VALUE KVM_DATASTORES.Accessible *EQ 'No'
```

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 condition specified by 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.

KVM_Datastore_Usage_High situation

Description
The data store is nearing or is at its defined capacity.

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

Formula
```
*IF *VALUE KVM_DATASTORES.Percent_Used *GT 90
```

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 condition specified by 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.

Events navigator item

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

KVM_Cluster_Critical_Event situation

Description
An error has occurred on the cluster.

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

Formula
```
*IF *VALUE KVM_EVENTS.Entity_Type *EQ 'Cluster' *AND *VALUE
KVM_EVENTS.Category *EQ 'error'
```

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
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Critical

Clearing conditions
The situation does not clear automatically.

KVM_Datastore_Critical_Event situation

Description
An error has occurred on the data store.

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

Formula

```
*IF *VALUE KVM_EVENTS.Entity_Type *EQ 'Datastore' *AND *VALUE
KVM_EVENTS.Category *EQ 'error'
```

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
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Critical

Clearing conditions
The situation does not clear automatically.

KVM_VM_Critical_Event situation

Description
An error has occurred on the virtual machine.

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

Formula

```
*IF *VALUE KVM_EVENTS.Entity_Type *EQ 'VirtualMachine' *AND *VALUE
KVM_EVENTS.Category *EQ 'error'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Critical

Clearing conditions
The situation does not clear automatically.

Monitored Servers navigator item

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

KVM_Collection_Error situation

Description
An ESX server is not responding to performance API queries.

The situation is evaluated for the table.

Formula
```
*IF *VALUE KVM_AGENT_EVENTS.Message *EQ 23 *AND *VALUE
KVM_AGENT_EVENTS.Severity *EQ Warning
```

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
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Warning

Clearing conditions
The situation does not clear automatically.

KVM_Collection_Time_Excessive situation

Description
A data collection is taking excessively long.

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

Formula
```
*IF *VALUE KVM_VCENTERS.Current_CU_Execution_Time *GT 600000
```

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
2 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_Connection_Failure situation

Description
A problem exists with the data source connection.

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

Formula

```
*IF *VALUE KVM_VCENTERS.Agent_Connection *EQ 0
```

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 condition specified by 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.

KVM_Host_System_Created situation

Description
A new ESX server was created.

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

Formula

```
*IF *VALUE KVM_TOPOLOGICAL_EVENTS.Entity_Type *EQ 'Host System' *AND
*VALUE KVM_TOPOLOGICAL_EVENTS.Event_Type *EQ 'Created'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Informational

Clearing conditions
The situation does not clear automatically.

KVM_Host_System_Created2 situation

Description
A new ESX server was created.

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

Formula
```
*IF *VALUE KVM_TOPOLOGICAL_EVENTS.Entity_Type *EQ 'Host System' *AND
*VALUE KVM_TOPOLOGICAL_EVENTS.Event_Type *EQ 'Created'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Informational

Clearing conditions
The situation does not clear automatically.

KVM_Host_System_Destroyed situation

Description
An ESX server was destroyed.

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

Formula
```
*IF *VALUE KVM_TOPOLOGICAL_EVENTS.Entity_Type *EQ 'Host System' *AND
*VALUE KVM_TOPOLOGICAL_EVENTS.Event_Type *EQ 'Destroyed'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Informational

Clearing conditions
The situation does not clear automatically.

KVM_Host_System_Destroyed2 situation

Description
An ESX server was destroyed.

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

Formula

```
*IF *VALUE KVM_TOPOLOGICAL_EVENTS.Entity_Type *EQ 'Host System' *AND
*VALUE KVM_TOPOLOGICAL_EVENTS.Event_Type *EQ 'Destroyed'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Informational

Clearing conditions
The situation does not clear automatically.

KVM_Inventory_Out_Of_Date situation

Description
The agent inventory is out of date.

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

Formula

```
*IF *VALUE KVM_VCENTERS.Inventory_Age *GT 180000
```

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 condition specified by 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.

KVM_Take_Action_Failure situation

Description
A problem occurred during a Take Action command.

The situation is evaluated for the table.

Formula

```
*IF *VALUE KVM_AGENT_EVENTS.Subsystem *EQ Task *AND *VALUE
KVM_AGENT_EVENTS.Severity *EQ Warning
```

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
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Warning

Clearing conditions
The situation does not clear automatically.

KVM_Virtual_Machine_Created situation

Description
A new virtual machine was created.

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

Formula

```
*IF *VALUE KVM_TOPOLOGICAL_EVENTS.Entity_Type *EQ 'Virtual Machine' *AND
*VALUE KVM_TOPOLOGICAL_EVENTS.Event_Type *EQ 'Created'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Informational

Clearing conditions
The situation does not clear automatically.

KVM_Virtual_Machine_Created2 situation

Description
A new virtual machine was created.

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

Formula

```
*IF *VALUE KVM_TOPOLOGICAL_EVENTS.Entity_Type *EQ 'Virtual Machine' *AND
*VALUE KVM_TOPOLOGICAL_EVENTS.Event_Type *EQ 'Created'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Informational

Clearing conditions
The situation does not clear automatically.

KVM_Virtual_Machine_Destroyed situation

Description
A virtual machine was destroyed.

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

Formula

```
*IF *VALUE KVM_TOPOLOGICAL_EVENTS.Entity_Type *EQ 'Virtual Machine' *AND
*VALUE KVM_TOPOLOGICAL_EVENTS.Event_Type *EQ 'Destroyed'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Informational

Clearing conditions
The situation does not clear automatically.

KVM_Virtual_Machine_Destroyed2 situation

Description
A virtual machine was destroyed.

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

Formula

```
*IF *VALUE KVM_TOPOLOGICAL_EVENTS.Entity_Type *EQ 'Virtual Machine' *AND
*VALUE KVM_TOPOLOGICAL_EVENTS.Event_Type *EQ 'Destroyed'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Informational

Clearing conditions
The situation does not clear automatically.

KVM_Virtual_Machine_Relocated situation

Description
A virtual machine was relocated.

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

Formula

```
*IF *VALUE KVM_TOPOLOGICAL_EVENTS.Entity_Type *EQ 'Virtual Machine' *AND
*VALUE KVM_TOPOLOGICAL_EVENTS.Event_Type *EQ 'Relocated'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Informational

Clearing conditions
The situation does not clear automatically.

KVM_Virtual_Machine_Relocated2 situation

Description
A virtual machine was relocated.

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

Formula

```
*IF *VALUE KVM_TOPOLOGICAL_EVENTS.Entity_Type *EQ 'Virtual Machine' *AND
*VALUE KVM_TOPOLOGICAL_EVENTS.Event_Type *EQ 'Relocated'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Informational

Clearing conditions
The situation does not clear automatically.

Networks navigator item

No predefined situations are included for this navigator item.

VMware VI subnode

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

VMware VI navigator item

No predefined situations are included for this navigator item.

CPU navigator item

KVM_VM_CPU_Ready_High situation

Description
The CPU percent ready is high.

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

Formula
```
*IF *VALUE KVM_VM_CPU.Percent_Rdy *GT 15
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_VM_CPU_Util_High situation

Description
The CPU utilization is high.

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

Formula
```
*IF *VALUE KVM_VM_CPU.Utilization *GT 90
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

Disk navigator item

KVM_Server_Disk_Reads_High situation

Description
The disk read activity is high.

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

Formula
```
*IF *VALUE KVM_SERVER_DISK.Read *GT 5000
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_Server_Disk_Writes_High situation

Description
The disk write activity is high.

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

Formula
```
*IF *VALUE KVM_SERVER_DISK.Write *GT 5000
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_VM_Disk_Free_Low situation

Description
The virtual machine disk partition free space is low.

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

Formula

```
*IF *VALUE KVM_VM_PARTITION.Percent_Free *GE 0 *AND *VALUE
KVM_VM_PARTITION.Percent_Free *LT 10
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

ESX Server navigator item

KVM_ESX_Server_Disconnected situation

Description
An ESX Server is not connected.

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

Formula

```
*IF *VALUE KVM_SERVER.Connection_State *NE 'Unavailable' *AND *VALUE
KVM_SERVER.Connection_State *NE 'connected' *AND *VALUE
KVM_SERVER.Maintenance_Mode *NE Yes
```

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
1 minute 30 seconds

Situation persistence
The number of times the condition specified by 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.

KVM_Host_Server_Bad_Status situation

Description
The status of the host server is not green.

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

Formula

```
*IF *VALUE KVM_SERVER.Overall_Status *NE 'Unavailable' *AND *VALUE
KVM_SERVER.Overall_Status *NE 'green'
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_Server_CPU_Util_High situation

Description
The CPU utilization is high.

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

Formula

```
*IF *VALUE KVM_SERVER.Overall_CPU_Util *GE 90
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_Server_Critical_Event situation

Description
An error has occurred on the ESX server.

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

Formula

```
*IF *VALUE KVM_SUBNODE_EVENTS.Entity_Type *EQ 'HostSystem' *AND *VALUE
KVM_SUBNODE_EVENTS.Category *EQ 'error'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Critical

Clearing conditions
The situation does not clear automatically.

KVM_Server_Datastore_Free_Low situation

Description
The data store free space is low.

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

Formula

```
*IF *VALUE KVM_SERVER_DATASTORE.Percent_Free *GE 0 *AND *VALUE
KVM_SERVER_DATASTORE.Percent_Free *LT 10
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_Server_HBA_Fault situation

Description
An ESX server host bus adapter has a fault.

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

Formula
```
*IF *VALUE KVM_SERVER_HBA.Status *EQ 'fault'
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_Server_Memory_Util_High situation

Description
The memory utilization is high.

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

Formula
```
*IF *VALUE KVM_SERVER.Overall_Memory_Util *GE 90
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_Server_VMotion_Event situation

Description
A VMotion event has been detected.

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

Formula
```
*IF *VALUE KVM_SUBNODE_EVENTS.Event_Type *EQ 'VmMigratedEvent' *OR *VALUE
KVM_SUBNODE_EVENTS.Event_Type *EQ 'DrsVmMigratedEvent'
```

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
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Informational

Clearing conditions
The situation does not clear automatically.

KVM_Server_VM_Critical_Event situation

Description
An error has occurred on the virtual machine.

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

Formula
```
*IF *VALUE KVM_SUBNODE_EVENTS.Entity_Type *EQ 'VirtualMachine' *AND
*VALUE KVM_SUBNODE_EVENTS.Category *EQ 'error'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
None. Data is analyzed when it becomes available.

Situation persistence
Not Applicable

Error conditions
Critical

Clearing conditions
The situation does not clear automatically.

Memory navigator item

KVM_VM_Guest_Memory_Util_High situation

Description
The virtual machine guest memory usage is high.

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

Formula
```
*IF *VALUE KVM_VM_MEMORY.Guest_Util *GT 90
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_VM_Host_Memory_Util_High situation

Description
The virtual machine host memory usage is high.

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

Formula
```
*IF *VALUE KVM_VM_MEMORY.Host_Util *GT 90
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

Network navigator item

KVM_Server_NIC_Down situation

Description
The host NIC adapter is not operational.

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

Formula

```
*IF *VALUE KVM_SERVER_NETWORK.Status *EQ 'down' *AND *VALUE
KVM_SERVER_NETWORK.Virtual_Switch *NE 'Unavailable'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_Server_Receive_Rate_High situation

Description
The receive rate is high for the server.

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

Formula

```
*IF *VALUE KVM_SERVER_NETWORK.Received *GT 5000
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_Server_Transmit_Rate_High situation

Description
The transmit rate is high for the server.

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

Formula
```
*IF *VALUE KVM_SERVER_NETWORK.Transmitted *GT 5000
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_VM_Receive_Rate_High situation

Description
The receive rate is high for the virtual machine.

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

Formula
```
*IF *VALUE KVM_VM_NETWORK.Received *GT 5000
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_VM_Transmit_Rate_High situation

Description
The transmit rate is high for the virtual machine.

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

Formula
```
*IF *VALUE KVM_VM_NETWORK.Transmitted *GT 5000
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
15 minutes

Situation persistence
The number of times the condition specified by 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.

Resource Pools navigator item

KVM_Resource_Pool_CPU_High situation

Description
The CPU utilization is high.

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

Formula
```
*IF *VALUE KVM_RESOURCE_POOL_CPU.Percent_Overall_Usage *GE 90
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_Resource_Pool_Memory_High situation

Description
The memory utilization is high.

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

Formula
```
*IF *VALUE KVM_RESOURCE_POOL_MEMORY.Percent_Overall_Usage *GE 90
```

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
15 minutes

Situation persistence
The number of times the condition specified by 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.

Virtual Machines navigator item

KVM_Snapshots_High situation

Description
The number of snapshots is high.

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

Formula
```
*IF *VALUE KVM_VIRTUAL_MACHINES.Number_Of_Snapshots *GE 32
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_VM_Bad_Status situation

Description
The status of the virtual machine is not green.

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

Formula
```
*IF *VALUE KVM_VIRTUAL_MACHINES.Overall_Status *NE 'Unavailable' *AND
*VALUE KVM_VIRTUAL_MACHINES.Overall_Status *NE 'green'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
15 minutes

Situation persistence
The number of times the condition specified by 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.

KVM_VM_Powered_Off situation

Description
The virtual machine is powered off.

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

Formula
```
*IF *VALUE KVM_VIRTUAL_MACHINES.Power_Status *EQ 'poweredOff'
```

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

Distribution
This situation is available for distribution.

Run at startup
No

Sampling interval
15 minutes

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

Error conditions
Informational

Clearing conditions
The situation clears when the condition becomes false.

Chapter 4. Take Action commands

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

About Take Action commands

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

In advanced automation, policies are used to take actions, schedule work, and automate manual tasks. A policy comprises a series of automated steps called activities that are connected to create a workflow. After an activity is completed, the Tivoli Enterprise Portal receives return-code feedback, and advanced automation logic responds with subsequent activities that are prescribed by the feedback.

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

Additional information about Take Action commands

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

For a list of the Take Action commands for this monitoring agent and a description of each command, see “Predefined Take Action commands” on page 201 and the information for each individual command.

Predefined Take Action commands

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

This monitoring agent contains the following Take Action commands:

- PowerOffVM
- PowerOnVM

Take Action command descriptions

Each Take Action command description provides information you can use to decide whether to run the Take Action command or whether to include the Take Action command in a situation or a policy.

The descriptions of the Take Action commands provide the following information:

Description
Actions the command performs on the system to which it is sent, and the permissions required for the Take Action command to function.

Return codes
Information that the Take Action command returns.

PowerOffVM action

This action attempts to power off a virtual machine. Two parameters are required for this action: the host name of the ESX server and the name of the virtual machine (the display name, not the virtual machine host name).

System command

To include the Take Action command in a situation or workflow policy, use the following syntax for the system command:

```
POWEROFFVM \
   [KVM_VIRTUAL_MACHINES.VM_Name]
```

You can use attribute substitution to supply the Take Action command arguments from the situation, for example:

```
POWEROFFVM \
[&{KVM_VIRTUAL_MACHINES.VM_Name}]
```

You can also use attribute substitution in a workflow policy though the format is slightly different:

```
POWEROFFVM \
[&WaitOnSituation:KVM_VIRTUAL_MACHINES.VM_Name]
```

Authorization

No authorization information provided.

Command arguments

- **Name:** KVM_VIRTUAL_MACHINES.VM_Name
  - **Description:** Name of the virtual machine to be powered off.
  - **Default:** ""

Destination systems

_EnDDESTINATIONS_NONE_OR_LIST_EnD

Return codes

- Return Code: 8
  - Return Code Type: TIMED_OUT
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM1019
  - Message: WARNING::NO MESSAGE FOUND FOR THIS RETURN CODE!!!!!
- Return Code: 12
  - Return Code Type: INSUFFICIENT_USER_AUTHORITY
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM1020
  - Message: WARNING::NO MESSAGE FOUND FOR THIS RETURN CODE!!!!!
- Return Code: 0
  - Return Code Type: OK
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)

  - Message ID: KVM5004I
  - Message: The request to power off the virtual machine was sent successfully.
- Return Code: 1
  - Return Code Type: NOT_RUNNING
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5005I
  - Message: The virtual machine is powered off.
- Return Code: 2
  - Return Code Type: GENERAL_ERROR
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5006E
  - Message: Could not perform the requested power off action.
- Return Code: 3
  - Return Code Type: GENERAL_ERROR
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5007E
  - Message: The ESX server name specified is invalid or could not be found.
- Return Code: 4
  - Return Code Type: GENERAL_ERROR
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5008E
  - Message: One of the required parameters for this action was not specified.
- Return Code: 5
  - Return Code Type: GENERAL_ERROR
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5009E
  - Message: An unknown action was specified for this request.
- Return Code: 6
  - Return Code Type: GENERAL_ERROR
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5045E
  - Message: The specified virtual machine was not found.

Usage

No authorization information provided.

PowerOnVM action

This action attempts to power on a virtual machine. Two parameters are required for this action: the host name of the ESX server and the name of the virtual machine (the display name, not the virtual machine host name).

System command

To include the Take Action command in a situation or workflow policy, use the following syntax for the system command:

```
POWERONVM \
   [KVM_VIRTUAL_MACHINES.VM_Name]
   [KVM_VIRTUAL_MACHINES.VM_Server_Name]
```

You can use attribute substitution to supply the Take Action command arguments from the situation, for example:

```
POWERONVM \
[&{KVM_VIRTUAL_MACHINES.VM_Name}] \
[&{KVM_VIRTUAL_MACHINES.VM_Server_Name}]
```

You can also use attribute substitution in a workflow policy though the format is slightly different:

```
POWERONVM \
[&WaitOnSituation:KVM_VIRTUAL_MACHINES.VM_Name] \
[&WaitOnSituation:KVM_VIRTUAL_MACHINES.VM_Server_Name]
```

Authorization

No authorization information provided.

Command arguments

- **Name:** KVM_VIRTUAL_MACHINES.VM_Name
  - **Description:** Name of the virtual machine to be powered on.
  - **Default:** ""
- **Name:** KVM_VIRTUAL_MACHINES.VM_Server_Name
  - **Description:** Name of the target virtual machine server.
  - **Default:** ""

Destination systems

_EnDDESTINATIONS_NONE_OR_LIST_EnD

Return codes

- Return Code: 8
  - Return Code Type: TIMED_OUT
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM1019
  - Message: WARNING::NO MESSAGE FOUND FOR THIS RETURN CODE!!!!!
- Return Code: 12
  - Return Code Type: INSUFFICIENT_USER_AUTHORITY

  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM1020
  - Message: WARNING::NO MESSAGE FOUND FOR THIS RETURN CODE!!!!!
- Return Code: 0
  - Return Code Type: OK
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5001I
  - Message: The request to power on the virtual machine was sent successfully.
- Return Code: 1
  - Return Code Type: ALREADY_RUNNING
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5002I
  - Message: The virtual machine is already powered on.
- Return Code: 2
  - Return Code Type: GENERAL_ERROR
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5003E
  - Message: Could not perform the requested power on action.
- Return Code: 3
  - Return Code Type: GENERAL_ERROR
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5007E
  - Message: The ESX server name specified is invalid or could not be found.
- Return Code: 4
  - Return Code Type: GENERAL_ERROR
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5008E
  - Message: One of the required parameters for this action was not specified.
- Return Code: 5
  - Return Code Type: GENERAL_ERROR
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)
  - Message ID: KVM5009E
  - Message: An unknown action was specified for this request.
- Return Code: 6
  - Return Code Type: GENERAL_ERROR
  - Operating systems: Linux (32-bit x86), Linux (64-bit IBM Z), Linux (64-bit x86), Windows (32-bit), Windows (64-bit)

- Message ID: KVM5045E
- Message: The specified virtual machine was not found.

Usage

No authorization information provided.

Chapter 5. Policies

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

A *policy* is a set of automated system processes that can take actions, schedule work for users, or automate manual tasks. You use the Workflow Editor to design policies. You control the order in which the policy executes a series of automated steps, which are also called *activities*. Policies are connected to create a workflow. After an activity is completed, the Tivoli Enterprise Portal receives return-code feedback, and advanced automation logic responds with subsequent activities prescribed by the feedback.

For more information about working with policies, see "Automation with policies" in the *Tivoli Enterprise Portal User's Guide*.

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

Predefined policies

This monitoring agent contains predefined workflow policies that interact with Tivoli Application Dependency Discovery Manager systems to keep the VMware topology up-to-date between scheduled discoveries performed by the Tivoli Application Dependency Discovery Manager sensors.

All these predefined policies are, by default, configured to send requests to the Tivoli Application Dependency Discovery Manager system identified by the name of VMWARE-TADDM in IBM Tivoli Monitoring.

In order to create the VMWARE-TADDM Tivoli Application Dependency Discovery Manager system in IBM Tivoli Monitoring, see the instructions listed in the section on initialization of Tivoli Application Dependency Discovery Manager policies in the Tivoli Enterprise Portal User’s Guide.

After the VMWARE-TADDM Tivoli Application Dependency Discovery Manager system is created in IBM Tivoli Monitoring, complete the following steps to enable the predefined policies to run:

1. Click the **Workflow Editor** icon.
2. Select a VMware Workflow policy and select the **Auto start** check box.
3. Ensure the policy is configured.
4. To save your changes, click **OK** or **Apply**.

This monitoring agent contains the following policies:

- KVM_VM_Created
- KVM_VM_Deleted
- KVM_VM_Relocated
- KVM_VMotion

KVM_VM_Created

This policy sends a create request to Tivoli Application Dependency Discovery Manager when a new virtual machine is created.

The create request is sent so that the corresponding virtual machine CDM object is created in the Tivoli Application Dependency Discovery Manager database. This policy is triggered by the KVM_Virtual_Machine_Created situation.

This policy includes two workflow activities:

On Demand Report activity
Used to collect additional information about the virtual machine being created that is not present in the situation processed, for example, the virtual machine name.

Send a Tivoli Application Dependency Discovery Manager Update activity
Used to send a create update to Tivoli Application Dependency Discovery Manager. The payload consists of the data contained in the situation processed and the data returned by the On Demand Report.

KVM_VM_Deleted

This policy sends a delete request to Tivoli Application Dependency Discovery Manager when a new virtual machine is deleted.

The delete request is sent so that the corresponding virtual machine CDM object is deleted from the Tivoli Application Dependency Discovery Manager database. This policy is triggered by the KVM_Virtual_Machine_Deleted situation.

This policy consists of a single workflow activity:

Send a Tivoli Application Dependency Discovery Manager Update activity
Used to send a delete update to Tivoli Application Dependency Discovery Manager. The payload consists of the data contained in the situation processed.

KVM_VM_Relocated

This policy sends a move request to Tivoli Application Dependency Discovery Manager when a virtual machine disk storage is moved.

The move request is sent so that the virtualizes relationship of the corresponding virtual machine CDM object is updated in the Tivoli Application Dependency Discovery Manager database. This policy is triggered by the KVM_Virtual_Machine_Relocated situation.

This policy consists of a single workflow activity:

Send a Tivoli Application Dependency Discovery Manager Update activity
Used to send a move update to Tivoli Application Dependency Discovery Manager. The payload consists of the data contained in the situation processed.

KVM_VMotion

This policy sends a move request to Tivoli Application Dependency Discovery Manager when a virtual machine is moved to execute somewhere else.

The move request is sent so that the virtualizes relationship of the corresponding virtual machine CDM object is updated in the Tivoli Application Dependency Discovery Manager database. This policy is triggered by the KVM_Server_VMotion situation.

This policy consists of a single workflow activity:

Send a Tivoli Application Dependency Discovery Manager Update activity
Used to send a move update to Tivoli Application Dependency Discovery Manager. The payload consists of the data contained in the situation processed.

Chapter 6. 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 VMware VI 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 KVM_Base and is defined in the `kvm.baroc` (version 7.30.03) file. The KVM_Base event class can be used for generic rules processing for any event from the IBM Tivoli Monitoring for Virtual Environments Agent for VMware VI.

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

- cancelable: INTEGER
- cancelable_enum: STRING
- initiated_by: STRING
- initiated_by_enum: STRING
- kvm_status: STRING
- kvm_status_enum: STRING
- name: STRING
- name_enum: STRING
- node: STRING
- queue_time: STRING
- queue_time_enum: STRING
- source_hostname: STRING
- source_hostname_enum: STRING
- start_time: STRING
- start_time_enum: STRING
- target_entity: STRING
- target_entity_enum: STRING
- target_entity_type: STRING

- target_entity_type_enum: STRING
- timestamp: STRING

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

- kvm_severity: INTEGER
- kvm_severity_enum: STRING
- kvm_source: STRING
- managed_system: STRING
- message: INTEGER
- message_enum: STRING
- node: STRING
- subsystem: INTEGER
- subsystem_enum: STRING
- timestamp: STRING

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

- cluster: STRING
- cluster_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- drs_type: STRING
- drs_type_enum: STRING
- fault_message: STRING
- fault_message_enum: STRING
- fault_name: STRING
- fault_name_enum: STRING
- ft_virtual_machine: STRING
- ft_virtual_machine_enum: STRING
- kvm_source: STRING
- kvm_source_enum: STRING
- node: STRING
- reason: STRING
- reason_enum: STRING
- source_hostname: STRING
- source_hostname_enum: STRING
- target_hostname: STRING
- target_hostname_enum: STRING
- timestamp: STRING
- virtual_machine: STRING

- virtual_machine_enum: STRING

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

- accessible: INTEGER
- accessible_enum: STRING
- capacity: INTEGER
- capacity_enum: STRING
- cluster: STRING
- cluster_enum: STRING
- cluster_moref: STRING
- cluster_moref_enum: STRING
- connected_hosts: INTEGER
- connected_hosts_enum: STRING
- connected_vms: INTEGER
- connected_vms_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- datastore: STRING
- datastore_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- msn: STRING
- msn_enum: STRING
- node: STRING
- nodeid: STRING
- overall_status: STRING
- overall_status_enum: STRING
- percent_used: INTEGER
- percent_used_enum: STRING
- remote_host_address: STRING
- remote_host_address_enum: STRING
- remote_path: STRING
- remote_path_enum: STRING
- timestamp: STRING
- type: STRING
- type_enum: STRING

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

- cluster_name: STRING

- cluster_name_enum: STRING
- cpu_usage: INTEGER
- cpu_usage_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- max_cpu_usage: INTEGER
- max_cpu_usage_enum: STRING
- max_memory_usage: INTEGER
- max_memory_usage_enum: STRING
- memory_usage: INTEGER
- memory_usage_enum: STRING
- node: STRING
- nodeid: STRING
- nodetype: STRING
- overall_status: STRING
- overall_status_enum: STRING
- percent_cpu_usage: INTEGER
- percent_cpu_usage_enum: STRING
- percent_memory_usage: INTEGER
- percent_memory_usage_enum: STRING
- pool_name: STRING
- pool_name_enum: STRING
- timestamp: STRING

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

- cluster_name: STRING
- cluster_name_enum: STRING
- cpu_effective_contribution: INTEGER
- cpu_effective_contribution_enum: STRING
- cpu_effective_utilization: INTEGER
- cpu_effective_utilization_enum: STRING
- cpu_total_contribution: INTEGER
- cpu_total_contribution_enum: STRING
- cpu_total_utilization: INTEGER
- cpu_total_utilization_enum: STRING

- datacenter: STRING
- datacenter_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- include_data_in_summarization_3: INTEGER
- include_data_in_summarization_3_enum: STRING
- memory_effective_utilization: INTEGER
- memory_effective_utilization_enum: STRING
- memory_total_utilization: INTEGER
- memory_total_utilization_enum: STRING
- mem_effective_contribution: INTEGER
- mem_effective_contribution_enum: STRING
- mem_total_contribution: INTEGER
- mem_total_contribution_enum: STRING
- msn_name: STRING
- msn_name_enum: STRING
- node: STRING
- nodeid: STRING
- overall_status: STRING
- overall_status_enum: STRING
- server_cpu_utilization: INTEGER
- server_cpu_utilization_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- server_memory_utilization: INTEGER
- server_memory_utilization_enum: STRING
- timestamp: STRING

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

- cluster_name: STRING
- cluster_name_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- destroy_with_parent: INTEGER

- destroy_with_parent_enum: STRING
- moref: STRING
- moref_enum: STRING
- node: STRING
- nodeid: STRING
- start_action: STRING
- start_action_enum: STRING
- start_delay: INTEGER
- start_delay_enum: STRING
- start_order: INTEGER
- start_order_enum: STRING
- stop_action: STRING
- stop_action_enum: STRING
- stop_delay: INTEGER
- stop_delay_enum: STRING
- timestamp: STRING
- virtual_app_name: STRING
- virtual_app_name_enum: STRING
- virtual_machine_name: STRING
- virtual_machine_name_enum: STRING
- waiting_for_guest: INTEGER
- waiting_for_guest_enum: STRING

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

- cluster_moref: STRING
- cluster_moref_enum: STRING
- cluster_name: STRING
- cluster_name_enum: STRING
- cpu_utilization: INTEGER
- cpu_utilization_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- memory_utilization: INTEGER
- memory_utilization_enum: STRING
- msn_name: STRING
- msn_name_enum: STRING

- node: STRING
- nodeid: STRING
- overall_status: STRING
- overall_status_enum: STRING
- timestamp: STRING
- vm_name: STRING
- vm_name_enum: STRING

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

- cluster_moref: STRING
- cluster_moref_enum: STRING
- cluster_name: STRING
- cluster_name_enum: STRING
- cpu_00_10: INTEGER
- cpu_00_10_enum: STRING
- cpu_10_20: INTEGER
- cpu_10_20_enum: STRING
- cpu_20_30: INTEGER
- cpu_20_30_enum: STRING
- cpu_30_40: INTEGER
- cpu_30_40_enum: STRING
- cpu_40_50: INTEGER
- cpu_40_50_enum: STRING
- cpu_50_60: INTEGER
- cpu_50_60_enum: STRING
- cpu_60_70: INTEGER
- cpu_60_70_enum: STRING
- cpu_70_80: INTEGER
- cpu_70_80_enum: STRING
- cpu_80_90: INTEGER
- cpu_80_90_enum: STRING
- cpu_90_100: INTEGER
- cpu_90_100_enum: STRING
- cpu_utilization: REAL
- cpu_utilization_enum: STRING
- current_evc_mode: STRING
- current_evc_mode_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING

- datacenter_moref: STRING
- datacenter_moref_enum: STRING
- datastores_total_free_space: INTEGER
- datastores_total_free_space_enum: STRING
- datastores_total_space: INTEGER
- datastores_total_space_enum: STRING
- drs_enabled: INTEGER
- drs_enabled_enum: STRING
- effective_cpu: REAL
- effective_cpu_enum: STRING
- effective_memory: REAL
- effective_memory_enum: STRING
- effective_servers: INTEGER
- effective_servers_enum: STRING
- ha_enabled: INTEGER
- ha_enabled_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_10: INTEGER
- include_data_in_summarization_10_enum: STRING
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- include_data_in_summarization_3: INTEGER
- include_data_in_summarization_3_enum: STRING
- include_data_in_summarization_4: INTEGER
- include_data_in_summarization_4_enum: STRING
- include_data_in_summarization_5: INTEGER
- include_data_in_summarization_5_enum: STRING
- include_data_in_summarization_6: INTEGER
- include_data_in_summarization_6_enum: STRING
- include_data_in_summarization_7: INTEGER
- include_data_in_summarization_7_enum: STRING
- include_data_in_summarization_8: INTEGER
- include_data_in_summarization_8_enum: STRING
- include_data_in_summarization_9: INTEGER
- include_data_in_summarization_9_enum: STRING
- memory_00_10: INTEGER

- memory_00_10_enum: STRING
- memory_10_20: INTEGER
- memory_10_20_enum: STRING
- memory_20_30: INTEGER
- memory_20_30_enum: STRING
- memory_30_40: INTEGER
- memory_30_40_enum: STRING
- memory_40_50: INTEGER
- memory_40_50_enum: STRING
- memory_50_60: INTEGER
- memory_50_60_enum: STRING
- memory_60_70: INTEGER
- memory_60_70_enum: STRING
- memory_70_80: INTEGER
- memory_70_80_enum: STRING
- memory_80_90: INTEGER
- memory_80_90_enum: STRING
- memory_90_100: INTEGER
- memory_90_100_enum: STRING
- memory_utilization: REAL
- memory_utilization_enum: STRING
- node: STRING
- nodeid: STRING
- number_cpus: INTEGER
- number_cpus_enum: STRING
- number_servers: INTEGER
- number_servers_enum: STRING
- number_vmotions: INTEGER
- number_vmotions_enum: STRING
- number_vms: INTEGER
- number_vms_enum: STRING
- number_vms_on: INTEGER
- number_vms_on_enum: STRING
- overall_status: STRING
- overall_status_enum: STRING
- percent_datastore_usage: REAL
- percent_datastore_usage_enum: STRING
- percent_effective_cpu: INTEGER
- percent_effective_cpu_enum: STRING

- percent_effective_memory: INTEGER
- percent_effective_memory_enum: STRING
- percent_effective_servers: INTEGER
- percent_effective_servers_enum: STRING
- physical_nics: INTEGER
- physical_nics_down: INTEGER
- physical_nics_down_enum: STRING
- physical_nics_enum: STRING
- servers_in_maintenance_mode: INTEGER
- servers_in_maintenance_mode_enum: STRING
- timestamp: STRING
- total_cpu: REAL
- total_cpu_enum: STRING
- total_memory: REAL
- total_memory_enum: STRING
- total_vm_configured_memory: REAL
- total_vm_configured_memory_enum: STRING
- total_vm_provisioned_space: REAL
- total_vm_provisioned_space_enum: STRING

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

- cpu_utilization: REAL
- cpu_utilization_enum: STRING
- datacenter: STRING
- effective_cpu: INTEGER
- effective_cpu_enum: STRING
- effective_memory: INTEGER
- effective_memory_enum: STRING
- effective_servers: INTEGER
- effective_servers_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- include_data_in_summarization_3: INTEGER
- include_data_in_summarization_3_enum: STRING
- memory_utilization: REAL

- memory_utilization_enum: STRING
- node: STRING
- nodeid: STRING
- overall_status: STRING
- overall_status_enum: STRING
- percent_effective_servers: REAL
- percent_effective_servers_enum: STRING
- timestamp: STRING
- total_cpu: INTEGER
- total_cpu_enum: STRING
- total_memory: INTEGER
- total_memory_enum: STRING
- total_servers: INTEGER
- total_servers_enum: STRING

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

- capacity: REAL
- capacity_enum: STRING
- capacity_used: REAL
- capacity_used_enum: STRING
- config_status: STRING
- config_status_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- datastore_cluster: STRING
- datastore_cluster_enum: STRING
- datastore_count: INTEGER
- datastore_count_enum: STRING
- default_intravm_affinity: INTEGER
- default_intravm_affinity_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- io_load_balance_enabled: INTEGER
- io_load_balance_enabled_enum: STRING
- load_balance_interval: INTEGER
- load_balance_interval_enum: STRING
- node: STRING

- overall_status: STRING
- overall_status_enum: STRING
- percent_capacity_free: INTEGER
- percent_capacity_free_enum: STRING
- timestamp: STRING

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

- datacenter: STRING
- datacenter_enum: STRING
- datastore: STRING
- datastore_enum: STRING
- disk: STRING
- disk_enum: STRING
- host: STRING
- host_enum: STRING
- node: STRING
- nodeid: STRING
- timestamp: STRING

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

- connectiontype: STRING
- connecttonode: STRING
- datacenter: STRING
- datacenter_enum: STRING
- msn: STRING
- msn_enum: STRING
- node: STRING
- nodeid: STRING
- nodename: STRING
- nodestatus: STRING
- nodetype: STRING
- timestamp: STRING

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

- accessible: INTEGER
- accessible_enum: STRING
- capacity: INTEGER
- capacity_enum: STRING
- connected_clusters: INTEGER
- connected_clusters_enum: STRING

- connected_hosts: INTEGER
- connected_hosts_enum: STRING
- connected_vms: INTEGER
- connected_vms_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- datastore_cluster: STRING
- datastore_cluster_enum: STRING
- datastore_moref: STRING
- datastore_moref_enum: STRING
- free_space: INTEGER
- free_space_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- include_data_in_summarization_3: INTEGER
- include_data_in_summarization_3_enum: STRING
- include_data_in_summarization_4: INTEGER
- include_data_in_summarization_4_enum: STRING
- include_data_in_summarization_5: INTEGER
- include_data_in_summarization_5_enum: STRING
- maximum_file_size: INTEGER
- maximum_file_size_enum: STRING
- msn: STRING
- msn_enum: STRING
- name: STRING
- name_enum: STRING
- netapp_volume_name: STRING
- netapp_volume_name_enum: STRING
- node: STRING
- nodeid: STRING
- overall_status: STRING
- overall_status_enum: STRING
- overcommitted: INTEGER
- overcommitted_enum: STRING
- percent_free: INTEGER

- percent_free_enum: STRING
- percent_overcommitted: REAL
- percent_overcommitted_enum: STRING
- percent_snapshot_storage_consumed: REAL
- percent_snapshot_storage_consumed_enum: STRING
- percent_used: INTEGER
- percent_used_enum: STRING
- remote_host_address: STRING
- remote_host_address_enum: STRING
- remote_path: STRING
- remote_path_enum: STRING
- snapshot_storage_consumed: REAL
- snapshot_storage_consumed_enum: STRING
- timestamp: STRING
- total_io_kbps: INTEGER
- total_io_kbps_enum: STRING
- total_read_kbps: INTEGER
- total_read_kbps_enum: STRING
- total_write_kbps: INTEGER
- total_write_kbps_enum: STRING
- type: STRING
- type_enum: STRING
- url: STRING
- url_enum: STRING
- used_space: INTEGER
- used_space_enum: STRING

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

- directorport: STRING
- directorserver: STRING
- node: STRING
- timestamp: STRING
- usetepcredential: STRING

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

- blocked: STRING
- blocked_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING

- inbound_shaping_average_bandwidth: INTEGER
- inbound_shaping_average_bandwidth_enum: STRING
- inbound_shaping_burst_size: INTEGER
- inbound_shaping_burst_size_enum: STRING
- inbound_shaping_enabled: STRING
- inbound_shaping_enabled_enum: STRING
- inbound_shaping_peak_bandwidth: INTEGER
- inbound_shaping_peak_bandwidth_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- node: STRING
- outbound_shaping_average_bandwidth: INTEGER
- outbound_shaping_average_bandwidth_enum: STRING
- outbound_shaping_burst_size: INTEGER
- outbound_shaping_burst_size_enum: STRING
- outbound_shaping_enabled: STRING
- outbound_shaping_enabled_enum: STRING
- outbound_shaping_peak_bandwidth: INTEGER
- outbound_shaping_peak_bandwidth_enum: STRING
- overall_status: STRING
- overall_status_enum: STRING
- portgroup_name: STRING
- portgroup_name_enum: STRING
- switch_name: STRING
- switch_name_enum: STRING
- timestamp: STRING
- type: STRING
- type_enum: STRING
- vlan_id: INTEGER
- vlan_id_enum: STRING
- vlan_type: STRING
- vlan_type_enum: STRING

For events that are generated by situations in the Distributed Virtual Switch Health attribute group, events are sent by using the ITM_KVM_DISTRIBUTED_VIRTUAL_SWITCH_HEALTH event class. This event class contains the following slots:

- datacenter: STRING
- datacenter_enum: STRING

- dvs_teaming_status: STRING
- dvs_teaming_status_enum: STRING
- health_check_type: STRING
- health_check_type_enum: STRING
- host: STRING
- host_enum: STRING
- kvm_source: STRING
- kvm_source_enum: STRING
- mtu_mismatch: STRING
- mtu_mismatch_enum: STRING
- nic_name: STRING
- nic_name_enum: STRING
- node: STRING
- portgroup_name: STRING
- portgroup_name_enum: STRING
- summary: STRING
- summary_enum: STRING
- switch_name: STRING
- switch_name_enum: STRING
- timestamp: STRING
- uplink_key: STRING
- uplink_key_enum: STRING
- uplink_name: STRING
- uplink_name_enum: STRING

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

- datacenter: STRING
- datacenter_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- include_data_in_summarization_3: INTEGER
- include_data_in_summarization_3_enum: STRING
- max_number_of_ports: INTEGER
- max_number_of_ports_enum: STRING

- node: STRING
- number_of_hosts: INTEGER
- number_of_hosts_enum: STRING
- number_of_portgroups: INTEGER
- number_of_portgroups_enum: STRING
- number_of_ports: INTEGER
- number_of_ports_enum: STRING
- number_of_uplinks: INTEGER
- number_of_uplinks_enum: STRING
- number_of_vms: INTEGER
- number_of_vms_enum: STRING
- overall_status: STRING
- overall_status_enum: STRING
- received: INTEGER
- received_enum: STRING
- switch_name: STRING
- switch_name_enum: STRING
- timestamp: STRING
- transmitted: INTEGER
- transmitted_enum: STRING
- usage: INTEGER
- usage_enum: STRING

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

- component_state: STRING
- component_state_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- duplex: STRING
- duplex_enum: STRING
- host_system: STRING
- host_system_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING

- link_speed: INTEGER
- link_speed_enum: STRING
- link_status: STRING
- link_status_enum: STRING
- link_utilization: REAL
- link_utilization_enum: STRING
- nic: STRING
- nic_enum: STRING
- node: STRING
- overall_status: STRING
- overall_status_enum: STRING
- portgroup_name: STRING
- portgroup_name_enum: STRING
- received: INTEGER
- received_enum: STRING
- subnode_msn: STRING
- subnode_msn_enum: STRING
- switch_name: STRING
- switch_name_enum: STRING
- timestamp: STRING
- transmitted: INTEGER
- transmitted_enum: STRING
- uplink_name: STRING
- uplink_name_enum: STRING
- usage: INTEGER
- usage_enum: STRING

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

- average_collection_duration: REAL
- average_collection_duration_enum: STRING
- cache_hits: INTEGER
- cache_hit_percent: REAL
- cache_misses: INTEGER
- error_code: INTEGER
- error_code_enum: STRING
- intervals_skipped: INTEGER
- last_collection_duration: REAL
- last_collection_finished: STRING

- last_collection_finished_enum: STRING
- last_collection_start: STRING
- last_collection_start_enum: STRING
- node: STRING
- number_of_collections: INTEGER
- object_name: STRING
- object_status: INTEGER
- object_status_enum: STRING
- object_type: INTEGER
- object_type_enum: STRING
- query_name: STRING
- refresh_interval: INTEGER
- timestamp: STRING

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

- category: STRING
- category_enum: STRING
- compute_resource: STRING
- compute_resource_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- datastore: STRING
- datastore_enum: STRING
- datastore_uuid: STRING
- datastore_uuid_enum: STRING
- entity_type: STRING
- entity_type_enum: STRING
- event: STRING
- event_enum: STRING
- event_seq_number: INTEGER
- event_seq_number_enum: STRING
- event_text: STRING
- event_text_enum: STRING
- event_time: STRING
- event_time_enum: STRING
- event_type: STRING
- event_type_enum: STRING
- event_type_id: STRING
- event_type_id_enum: STRING

- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- node: STRING
- source_hostname: STRING
- source_hostname_enum: STRING
- timestamp: STRING
- userid: STRING
- userid_enum: STRING
- virtual_machine: STRING
- virtual_machine_enum: STRING
- virtual_machine_uuid: STRING
- virtual_machine_uuid_enum: STRING

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

- node: STRING
- subnode_affinity: STRING
- subnode_msn: STRING
- subnode_resource_name: STRING
- subnode_resource_name_enhanced: STRING
- subnode_type: STRING
- subnode_version: STRING
- timestamp: STRING

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

- datacenter: STRING
- datacenter_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- network: STRING
- network_enum: STRING
- node: STRING
- received: INTEGER
- received_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- subnode_msn: STRING
- subnode_msn_enum: STRING
- switch: STRING
- switch_enum: STRING

- timestamp: STRING
- transmitted: INTEGER
- transmitted_enum: STRING
- usage: INTEGER
- usage_enum: STRING

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

- datacenter: STRING
- datacenter_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- network: STRING
- network_enum: STRING
- node: STRING
- received: INTEGER
- received_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- subnode_msn: STRING
- subnode_msn_enum: STRING
- switch: STRING
- switch_enum: STRING
- timestamp: STRING
- transmitted: INTEGER
- transmitted_enum: STRING
- usage: INTEGER
- usage_enum: STRING
- virtual_machine: STRING
- virtual_machine_enum: STRING
- vm_nic: STRING
- vm_nic_enum: STRING

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

- datacenter: STRING
- datacenter_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER

- include_data_in_summarization_1_enum: STRING
- network: STRING
- network_enum: STRING
- node: STRING
- number_of_nics: INTEGER
- number_of_nics_enum: STRING
- received: INTEGER
- received_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- subnode_msn: STRING
- subnode_msn_enum: STRING
- switch: STRING
- switch_enum: STRING
- timestamp: STRING
- transmitted: INTEGER
- transmitted_enum: STRING
- usage: INTEGER
- usage_enum: STRING

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

- datacenter: STRING
- datacenter_enum: STRING
- distributed_switch: STRING
- distributed_switch_enum: STRING
- network: STRING
- network_enum: STRING
- network_type: STRING
- network_type_enum: STRING
- node: STRING
- number_of_hosts: INTEGER
- number_of_hosts_enum: STRING
- number_of_vms: INTEGER
- number_of_vms_enum: STRING
- overall_status: STRING
- overall_status_enum: STRING
- timestamp: STRING

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

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

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

- cpu_usage: INTEGER
- cpu_usage_enum: STRING
- expandable: INTEGER
- expandable_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- include_data_in_summarization_3: INTEGER
- include_data_in_summarization_3_enum: STRING
- limit: INTEGER
- limit_enum: STRING

- max_usage: INTEGER
- max_usage_enum: STRING
- node: STRING
- nodeid: STRING
- parent_name: STRING
- parent_name_enum: STRING
- percent_overall_usage: INTEGER
- percent_overall_usage_enum: STRING
- percent_reserved_vms: INTEGER
- percent_reserved_vms_enum: STRING
- pool_name: STRING
- pool_name_enum: STRING
- reservation: INTEGER
- reservation_enum: STRING
- reservation_used: INTEGER
- reservation_used_enum: STRING
- reservation_used_vm: INTEGER
- reservation_used_vm_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- shares: INTEGER
- shares_enum: STRING
- share_level: STRING
- share_level_enum: STRING
- timestamp: STRING
- unreserved: INTEGER
- unreserved_enum: STRING
- unreserved_vm: INTEGER
- unreserved_vm_enum: STRING

For events that are generated by situations in the Resource Pool General attribute group, events are sent by using the ITM_KVM_RESOURCE_POOL_GENERAL event class. This event class contains the following slots:

- cpu_usage: INTEGER
- cpu_usage_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER

- include_data_in_summarization_2_enum: STRING
- kvm_status: STRING
- kvm_status_enum: STRING
- memory_usage: INTEGER
- memory_usage_enum: STRING
- node: STRING
- nodeid: STRING
- number_child_pools: INTEGER
- number_child_pools_enum: STRING
- number_vms: INTEGER
- number_vms_enum: STRING
- number_vms_on: INTEGER
- number_vms_on_enum: STRING
- parent_name: STRING
- parent_name_enum: STRING
- pool_name: STRING
- pool_name_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- timestamp: STRING

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

- expandable: INTEGER
- expandable_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- include_data_in_summarization_3: INTEGER
- include_data_in_summarization_3_enum: STRING
- limit: INTEGER
- limit_enum: STRING
- max_usage: INTEGER
- max_usage_enum: STRING
- memory_usage: INTEGER
- memory_usage_enum: STRING

- node: STRING
- nodeid: STRING
- parent_name: STRING
- parent_name_enum: STRING
- percent_overall_usage: INTEGER
- percent_overall_usage_enum: STRING
- percent_reserved_vms: INTEGER
- percent_reserved_vms_enum: STRING
- pool_name: STRING
- pool_name_enum: STRING
- reservation: INTEGER
- reservation_enum: STRING
- reservation_used: INTEGER
- reservation_used_enum: STRING
- reservation_used_vm: INTEGER
- reservation_used_vm_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- shares: INTEGER
- shares_enum: STRING
- share_level: STRING
- share_level_enum: STRING
- timestamp: STRING
- unreserved: INTEGER
- unreserved_enum: STRING
- unreserved_vm: INTEGER
- unreserved_vm_enum: STRING

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

- avg_vm_cpu_percent_rdy: REAL
- avg_vm_cpu_percent_rdy_enum: STRING
- bios_date: STRING
- bios_date_enum: STRING
- build_number: STRING
- build_number_enum: STRING
- cluster: STRING
- cluster_enum: STRING
- connection_state: STRING
- connection_state_enum: STRING

- cpu_packages: INTEGER
- cpu_packages_enum: STRING
- current_evc_mode: STRING
- current_evc_mode_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- datacenter_moref: STRING
- datacenter_moref_enum: STRING
- datastore_space: INTEGER
- datastore_space_enum: STRING
- demand: INTEGER
- demand_enum: STRING
- energy_usage: INTEGER
- energy_usage_enum: STRING
- fully_qualified_name: STRING
- fully_qualified_name_enum: STRING
- hyperthreading_enabled: INTEGER
- hyperthreading_enabled_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- include_data_in_summarization_3: INTEGER
- include_data_in_summarization_3_enum: STRING
- include_data_in_summarization_4: INTEGER
- include_data_in_summarization_4_enum: STRING
- include_data_in_summarization_5: INTEGER
- include_data_in_summarization_5_enum: STRING
- include_data_in_summarization_6: INTEGER
- include_data_in_summarization_6_enum: STRING
- include_data_in_summarization_7: INTEGER
- include_data_in_summarization_7_enum: STRING
- include_data_in_summarization_8: INTEGER
- include_data_in_summarization_8_enum: STRING
- include_data_in_summarization_9: INTEGER
- include_data_in_summarization_9_enum: STRING
- ip_address: STRING

- ip_address_enum: STRING
- latency: REAL
- latency_enum: STRING
- maintenance_mode: INTEGER
- maintenance_mode_enum: STRING
- max_evc_mode: STRING
- max_evc_mode_enum: STRING
- nics: INTEGER
- nics_enum: STRING
- node: STRING
- nodeid: STRING
- number_vms: INTEGER
- number_vms_enum: STRING
- number_vms_on: INTEGER
- number_vms_on_enum: STRING
- overall_cpu_util: INTEGER
- overall_cpu_util_enum: STRING
- overall_memory_util: INTEGER
- overall_memory_util_enum: STRING
- overall_status: STRING
- overall_status_enum: STRING
- performance_error_pct: REAL
- performance_error_pct_enum: STRING
- performance_error_rate: INTEGER
- performance_error_rate_enum: STRING
- physical_cpus: INTEGER
- physical_cpus_enum: STRING
- physical_memory: INTEGER
- physical_memory_enum: STRING
- power_capacity: INTEGER
- power_capacity_enum: STRING
- power_state: STRING
- power_state_enum: STRING
- power_usage: INTEGER
- power_usage_enum: STRING
- processor_family: STRING
- processor_family_enum: STRING
- product: STRING
- product_enum: STRING

- serial_number: STRING
- serial_number_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- ssh_status: INTEGER
- ssh_status_enum: STRING
- storage_adapter_max_latency: INTEGER
- storage_adapter_max_latency_enum: STRING
- storage_path_max_latency: INTEGER
- storage_path_max_latency_enum: STRING
- system_model: STRING
- system_model_enum: STRING
- system_up_time: INTEGER
- system_up_time_enum: STRING
- system_vendor: STRING
- system_vendor_enum: STRING
- timestamp: STRING
- total_cpu_mhz: INTEGER
- total_cpu_mhz_enum: STRING
- total_vm_configured_memory: REAL
- total_vm_configured_memory_enum: STRING
- total_vm_provisioned_space: REAL
- total_vm_provisioned_space_enum: STRING
- used_cpu_mhz: INTEGER
- used_cpu_mhz_enum: STRING
- used_datastore: INTEGER
- used_datastore_enum: STRING
- uuid: STRING
- uuid_enum: STRING
- version: STRING
- version_enum: STRING
- vmotion_enabled: STRING
- vmotion_enabled_enum: STRING

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

- core_utilization: REAL
- core_utilization_enum: STRING
- cpu_number: INTEGER
- cpu_number_enum: STRING

- cpu_utilization: INTEGER
- cpu_utilization_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- node: STRING
- nodeid: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- timestamp: STRING

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

- agent_msn: STRING
- agent_msn_enum: STRING
- capacity: INTEGER
- capacity_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- datastore_moref: STRING
- datastore_moref_enum: STRING
- free_space: INTEGER
- free_space_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- maximum_file_size: INTEGER
- maximum_file_size_enum: STRING
- name: STRING
- name_enum: STRING
- node: STRING
- nodeid: STRING
- overall_status: STRING
- overall_status_enum: STRING
- percent_free: INTEGER
- percent_free_enum: STRING
- percent_used: INTEGER
- percent_used_enum: STRING

- read_latency: INTEGER
- read_latency_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- timestamp: STRING
- type: STRING
- type_enum: STRING
- used_space: INTEGER
- used_space_enum: STRING
- write_latency: INTEGER
- write_latency_enum: STRING

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

- backing_datastore: STRING
- backing_datastore_enum: STRING
- bus_resets: INTEGER
- bus_resets_enum: STRING
- commands: INTEGER
- commands_aborted: INTEGER
- commands_aborted_enum: STRING
- commands_enum: STRING
- device_latency: INTEGER
- device_latency_enum: STRING
- device_read_latency: INTEGER
- device_read_latency_enum: STRING
- device_total_latency: INTEGER
- device_total_latency_enum: STRING
- device_write_latency: INTEGER
- device_write_latency_enum: STRING
- disk_name: STRING
- disk_name_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- include_data_in_summarization_3: INTEGER
- include_data_in_summarization_3_enum: STRING

- include_data_in_summarization_4: INTEGER
- include_data_in_summarization_4_enum: STRING
- include_data_in_summarization_5: INTEGER
- include_data_in_summarization_5_enum: STRING
- include_data_in_summarization_6: INTEGER
- include_data_in_summarization_6_enum: STRING
- include_data_in_summarization_7: INTEGER
- include_data_in_summarization_7_enum: STRING
- kernel_latency: INTEGER
- kernel_latency_enum: STRING
- kernel_read_latency: INTEGER
- kernel_read_latency_enum: STRING
- kernel_total_latency: INTEGER
- kernel_total_latency_enum: STRING
- kernel_write_latency: INTEGER
- kernel_write_latency_enum: STRING
- node: STRING
- nodeid: STRING
- number_read: INTEGER
- number_read_enum: STRING
- number_write: INTEGER
- number_write_enum: STRING
- queue_latency: INTEGER
- queue_latency_enum: STRING
- queue_read_latency: INTEGER
- queue_read_latency_enum: STRING
- queue_total_latency: INTEGER
- queue_total_latency_enum: STRING
- queue_write_latency: INTEGER
- queue_write_latency_enum: STRING
- read: INTEGER
- read_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- timestamp: STRING
- total_latency: INTEGER
- total_latency_enum: STRING
- total_read_latency: INTEGER
- total_read_latency_enum: STRING

- total_write_latency: INTEGER
- total_write_latency_enum: STRING
- write: INTEGER
- write_enum: STRING

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

- bus: INTEGER
- bus_enum: STRING
- current_link_speed: INTEGER
- current_link_speed_enum: STRING
- device: STRING
- device_enum: STRING
- driver: STRING
- driver_enum: STRING
- hba_type: STRING
- hba_type_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- kvm_status: STRING
- kvm_status_enum: STRING
- max_link_speed: INTEGER
- max_link_speed_enum: STRING
- model: STRING
- model_enum: STRING
- node: STRING
- nodeid: STRING
- pci_id: STRING
- pci_id_enum: STRING
- read: INTEGER
- read_enum: STRING
- read_latency: INTEGER
- read_latency_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- speed: INTEGER
- speed_enum: STRING
- storage_adapter_throughput_usage: INTEGER

- storage_adapter_throughput_usage_enum: STRING
- timestamp: STRING
- write: INTEGER
- write_enum: STRING
- write_latency: INTEGER
- write_latency_enum: STRING

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

- node: STRING
- nodeid: STRING
- sensor_name: STRING
- sensor_name_enum: STRING
- sensor_status: STRING
- sensor_status_enum: STRING
- sensor_type: STRING
- sensor_type_enum: STRING
- sensor_units: STRING
- sensor_units_enum: STRING
- sensor_value: REAL
- sensor_value_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- timestamp: STRING

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

- active_memory: INTEGER
- active_memory_enum: STRING
- active_write: INTEGER
- active_write_enum: STRING
- balloon_used: INTEGER
- balloon_used_enum: STRING
- free_memory: INTEGER
- free_memory_enum: STRING
- granted_max_memory: INTEGER
- granted_max_memory_enum: STRING
- granted_memory: INTEGER
- granted_memory_enum: STRING
- granted_min_memory: INTEGER
- granted_min_memory_enum: STRING

- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- include_data_in_summarization_3: INTEGER
- include_data_in_summarization_3_enum: STRING
- include_data_in_summarization_4: INTEGER
- include_data_in_summarization_4_enum: STRING
- low_free_threshold: INTEGER
- low_free_threshold_enum: STRING
- memory_usage: INTEGER
- memory_usage_enum: STRING
- memory_utilization: INTEGER
- memory_utilization_enum: STRING
- node: STRING
- nodeid: STRING
- physical_memory: INTEGER
- physical_memory_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- service_console: INTEGER
- service_console_enum: STRING
- swap_in_rate: INTEGER
- swap_in_rate_enum: STRING
- swap_in_rate_host_cache: INTEGER
- swap_in_rate_host_cache_enum: STRING
- swap_out_rate: INTEGER
- swap_out_rate_enum: STRING
- swap_out_rate_host_cache: INTEGER
- swap_out_rate_host_cache_enum: STRING
- swap_total_rate: INTEGER
- swap_total_rate_enum: STRING
- swap_used: INTEGER
- swap_used_enum: STRING
- timestamp: STRING

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

- cluster: STRING
- cluster_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- duplex: STRING
- duplex_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- kvm_status: STRING
- kvm_status_enum: STRING
- link_speed: INTEGER
- link_speed_enum: STRING
- link_utilization: REAL
- link_utilization_enum: STRING
- nic_name: STRING
- nic_name_enum: STRING
- node: STRING
- nodeid: STRING
- physical_addr: STRING
- physical_addr_enum: STRING
- pkts_dropped: INTEGER
- pkts_dropped_enum: STRING
- pkts_received: INTEGER
- pkts_received_enum: STRING
- pkts_transmitted: INTEGER
- pkts_transmitted_enum: STRING
- received: INTEGER
- received_enum: STRING
- receive_pkts_dropped: INTEGER
- receive_pkts_dropped_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- timestamp: STRING
- transmitted: INTEGER
- transmitted_enum: STRING

- transmit_pkts_dropped: INTEGER
- transmit_pkts_dropped_enum: STRING
- usage: INTEGER
- usage_enum: STRING
- virtual_switch: STRING
- virtual_switch_enum: STRING

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

- broken_paths: INTEGER
- broken_paths_enum: STRING
- datastore: STRING
- datastore_enum: STRING
- disabled_paths: INTEGER
- disabled_paths_enum: STRING
- disk_name: STRING
- disk_name_enum: STRING
- node: STRING
- nodeid: STRING
- paths: INTEGER
- paths_enum: STRING
- path_selection_policy: STRING
- path_selection_policy_enum: STRING
- timestamp: STRING

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

- datacenter: STRING
- datacenter_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- network: STRING
- network_enum: STRING
- node: STRING
- number_of_nics: INTEGER
- number_of_nics_enum: STRING
- received: INTEGER
- received_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- switch: STRING

- switch_enum: STRING
- timestamp: STRING
- transmitted: INTEGER
- transmitted_enum: STRING
- usage: INTEGER
- usage_enum: STRING

For events that are generated by situations in the Server VM Datastore Utilization attribute group, events are sent by using the ITM_KVM_SERVER_VM_DATASTORE_UTILIZATION event class. This event class contains the following slots:

- committed: REAL
- committed_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- name: STRING
- name_enum: STRING
- node: STRING
- nodeid: STRING
- percent_committed: REAL
- percent_committed_enum: STRING
- provisioned: REAL
- provisioned_enum: STRING
- timestamp: STRING
- uncommitted: REAL
- uncommitted_enum: STRING
- unshared: REAL
- unshared_enum: STRING
- virtual_machine: STRING
- virtual_machine_enum: STRING
- vmnodeid: STRING

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

- category: STRING
- category_enum: STRING
- compute_resource: STRING
- compute_resource_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- entity_type: STRING

- entity_type_enum: STRING
- esx_server_uuid: STRING
- esx_server_uuid_enum: STRING
- event: STRING
- event_enum: STRING
- event_seq_number: INTEGER
- event_seq_number_enum: STRING
- event_text: STRING
- event_text_enum: STRING
- event_time: STRING
- event_time_enum: STRING
- event_type: STRING
- event_type_enum: STRING
- event_type_id: STRING
- event_type_id_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- node: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- timestamp: STRING
- userid: STRING
- userid_enum: STRING
- virtual_machine: STRING
- virtual_machine_enum: STRING
- virtual_machine_uuid: STRING
- virtual_machine_uuid_enum: STRING

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

- completed_time: STRING
- completed_time_enum: STRING
- error_message: STRING
- error_message_enum: STRING
- initiated_by: STRING
- initiated_by_enum: STRING
- kvm_status: STRING
- kvm_status_enum: STRING
- name: STRING
- name_enum: STRING

- node: STRING
- queue_time: STRING
- queue_time_enum: STRING
- source_hostname: STRING
- source_hostname_enum: STRING
- start_time: STRING
- start_time_enum: STRING
- target_entity: STRING
- target_entity_enum: STRING
- target_entity_type: STRING
- target_entity_type_enum: STRING
- timestamp: STRING

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

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

- timestamp: STRING

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

- datastore_uuid: STRING
- datastore_uuid_enum: STRING
- entity_type: STRING
- entity_type_enum: STRING
- event_type: STRING
- event_type_enum: STRING
- host_uuid: STRING
- host_uuid_enum: STRING
- msn: STRING
- msn_enum: STRING
- name: STRING
- name_enum: STRING
- node: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- timestamp: STRING
- vm_uuid: STRING
- vm_uuid_enum: STRING

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

- connectiontype: STRING
- connecttonode: STRING
- datacenter: STRING
- datacenter_enum: STRING
- msn: STRING
- msn_enum: STRING
- node: STRING
- nodeid: STRING
- nodename: STRING
- nodestatus: STRING
- nodetype: STRING
- timestamp: STRING

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

- affected_entity: STRING
- affected_entity_enum: STRING
- alarm_name: STRING

- alarm_name_enum: STRING
- alarm_status: STRING
- alarm_status_enum: STRING
- alarm_triggered_time: STRING
- alarm_triggered_time_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- description: STRING
- description_enum: STRING
- node: STRING
- timestamp: STRING
- triggered_entity: STRING
- triggered_entity_enum: STRING

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

- agent_connection: INTEGER
- agent_connection_enum: STRING
- average_cu_execution_time: REAL
- average_cu_execution_time_enum: STRING
- average_cu_queue_time: REAL
- average_cu_queue_time_enum: STRING
- collection_units: INTEGER
- collection_units_enum: STRING
- configured_address: STRING
- current_cu_execution_time: REAL
- current_cu_execution_time_enum: STRING
- current_cu_queue_time: REAL
- current_cu_queue_time_enum: STRING
- executing_collection_units: INTEGER
- executing_collection_units_enum: STRING
- fqdn: STRING
- fqdn_enum: STRING
- inventory_age: REAL
- inventory_age_enum: STRING
- ip_address: STRING
- ip_address_enum: STRING
- node: STRING
- queued_collection_units: INTEGER
- queued_collection_units_enum: STRING

- timestamp: STRING
- type: STRING
- type_enum: STRING
- web_services_port: INTEGER
- web_services_port_enum: STRING

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

- cluster: STRING
- cluster_enum: STRING
- connection_state: STRING
- connection_state_enum: STRING
- consolidationneeded: INTEGER
- consolidationneeded_enum: STRING
- cpu_limit: INTEGER
- cpu_limit_enum: STRING
- cpu_reservation: INTEGER
- cpu_reservation_enum: STRING
- cpu_shares: INTEGER
- cpu_shares_enum: STRING
- cpu_utilization: INTEGER
- cpu_utilization_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- fault_tolerance: INTEGER
- fault_tolerance_enum: STRING
- ft_instance_uuid: STRING
- ft_instance_uuid_enum: STRING
- guestos_msn: STRING
- guestos_msn_enum: STRING
- guestos_name: STRING
- guestos_name_enum: STRING
- guest_state: STRING
- guest_state_enum: STRING
- heartbeats: INTEGER
- heartbeats_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING

- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING
- include_data_in_summarization_3: INTEGER
- include_data_in_summarization_3_enum: STRING
- include_data_in_summarization_4: INTEGER
- include_data_in_summarization_4_enum: STRING
- include_data_in_summarization_5: INTEGER
- include_data_in_summarization_5_enum: STRING
- include_data_in_summarization_6: INTEGER
- include_data_in_summarization_6_enum: STRING
- instance_uuid: STRING
- instance_uuid_enum: STRING
- ip_address: STRING
- ip_address_enum: STRING
- kvm_hostname: STRING
- kvm_hostname_enum: STRING
- memory_limit: INTEGER
- memory_limit_enum: STRING
- memory_reservation: INTEGER
- memory_reservation_enum: STRING
- memory_shares: INTEGER
- memory_shares_enum: STRING
- memory_size: INTEGER
- memory_size_enum: STRING
- moref: STRING
- moref_enum: STRING
- node: STRING
- nodeid: STRING
- number_of_disks: INTEGER
- number_of_disks_enum: STRING
- number_of_nics: INTEGER
- number_of_nics_enum: STRING
- number_of_snapshots: INTEGER
- number_of_snapshots_enum: STRING
- num_cpus: INTEGER
- num_cpus_enum: STRING
- overall_status: STRING
- overall_status_enum: STRING
- power_status: STRING

- power_status_enum: STRING
- resource_pool: STRING
- resource_pool_enum: STRING
- snapshot_storage_consumed: INTEGER
- snapshot_storage_consumed_enum: STRING
- storage_drs_enable: INTEGER
- storage_drs_enable_enum: STRING
- template: INTEGER
- template_enum: STRING
- timestamp: STRING
- tools_status: STRING
- tools_status_enum: STRING
- up_time: INTEGER
- up_time_enum: STRING
- used_cpu_mhz: INTEGER
- used_cpu_mhz_enum: STRING
- uuid: STRING
- uuid_enum: STRING
- version: STRING
- version_enum: STRING
- vm_datastorepercent_utilization: REAL
- vm_datastorepercent_utilization_enum: STRING
- vm_name: STRING
- vm_name_enum: STRING
- vm_os_type: INTEGER
- vm_os_type_enum: STRING
- vm_percent_rdy: REAL
- vm_percent_rdy_enum: STRING
- vm_server_name: STRING
- vm_server_name_enum: STRING

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

- datacenter: STRING
- datacenter_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- node: STRING

- number_of_nics: INTEGER
- number_of_nics_enum: STRING
- received: INTEGER
- received_enum: STRING
- server_hostname: STRING
- server_hostname_enum: STRING
- subnode_msn: STRING
- subnode_msn_enum: STRING
- switch: STRING
- switch_enum: STRING
- timestamp: STRING
- transmitted: INTEGER
- transmitted_enum: STRING
- usage: INTEGER
- usage_enum: STRING

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

- cpu_number: INTEGER
- cpu_number_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- node: STRING
- nodeid: STRING
- percent_rdy: INTEGER
- percent_rdy_enum: STRING
- ready_time: INTEGER
- ready_time_enum: STRING
- sys_time: INTEGER
- sys_time_enum: STRING
- timestamp: STRING
- used_time: INTEGER
- used_time_enum: STRING
- user_time: INTEGER
- user_time_enum: STRING
- utilization: INTEGER
- utilization_enum: STRING
- vm_hostname: STRING
- vm_hostname_enum: STRING
- vm_name: STRING

- vm_name_cpu_number: STRING
- vm_name_cpu_number_enum: STRING
- vm_name_enum: STRING
- vm_os_type: INTEGER
- vm_os_type_enum: STRING
- vm_server_name: STRING
- vm_server_name_enum: STRING
- wait_time: INTEGER
- wait_time_enum: STRING

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

- committed: REAL
- committed_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- name: STRING
- name_enum: STRING
- node: STRING
- nodeid: STRING
- percent_committed: REAL
- percent_committed_enum: STRING
- provisioned: REAL
- provisioned_enum: STRING
- timestamp: STRING
- total_io_kbps: INTEGER
- total_io_kbps_enum: STRING
- total_read_kbps: INTEGER
- total_read_kbps_enum: STRING
- total_write_kbps: INTEGER
- total_write_kbps_enum: STRING
- uncommitted: REAL
- uncommitted_enum: STRING
- unshared: REAL
- unshared_enum: STRING

- uuid: STRING
- virtual_machine: STRING
- virtual_machine_enum: STRING

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

- access: STRING
- access_enum: STRING
- backing_datastore: STRING
- backing_datastore_enum: STRING
- capacity: INTEGER
- capacity_enum: STRING
- connected: STRING
- connected_enum: STRING
- description: STRING
- description_enum: STRING
- disk_shares: INTEGER
- disk_shares_enum: STRING
- node: STRING
- nodeid: STRING
- removable: STRING
- removable_enum: STRING
- timestamp: STRING
- vm_hostname: STRING
- vm_hostname_enum: STRING
- vm_name: STRING
- vm_name_enum: STRING
- vm_os_type: INTEGER
- vm_os_type_enum: STRING
- vm_server_name: STRING
- vm_server_name_enum: STRING

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

- disk_name: STRING
- disk_name_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- moref: STRING
- moref_enum: STRING
- node: STRING

- number_read: INTEGER
- number_read_enum: STRING
- number_write: INTEGER
- number_write_enum: STRING
- read: INTEGER
- read_enum: STRING
- timestamp: STRING
- virtual_machine: STRING
- virtual_machine_enum: STRING
- write: INTEGER
- write_enum: STRING

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

- active: INTEGER
- active_enum: STRING
- balloon_usage: INTEGER
- balloon_usage_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- granted: INTEGER
- granted_enum: STRING
- guest_free: INTEGER
- guest_free_enum: STRING
- guest_usage: INTEGER
- guest_usage_enum: STRING
- guest_util: INTEGER
- guest_util_enum: STRING
- host_free: INTEGER
- host_free_enum: STRING
- host_usage: INTEGER
- host_usage_enum: STRING
- host_util: INTEGER
- host_util_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- include_data_in_summarization_1: INTEGER
- include_data_in_summarization_1_enum: STRING
- include_data_in_summarization_2: INTEGER
- include_data_in_summarization_2_enum: STRING

- max_alloc: INTEGER
- max_alloc_enum: STRING
- min_alloc: INTEGER
- min_alloc_enum: STRING
- node: STRING
- nodeid: STRING
- shared: INTEGER
- shared_enum: STRING
- swap_in_rate: INTEGER
- swap_in_rate_enum: STRING
- swap_out_rate: INTEGER
- swap_out_rate_enum: STRING
- swap_to_file: INTEGER
- swap_to_file_enum: STRING
- timestamp: STRING
- total_size: INTEGER
- total_size_enum: STRING
- usage: REAL
- usage_enum: STRING
- vm_hostname: STRING
- vm_hostname_enum: STRING
- vm_name: STRING
- vm_name_enum: STRING
- vm_os_type: INTEGER
- vm_os_type_enum: STRING
- vm_server_name: STRING
- vm_server_name_enum: STRING

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

- cluster: STRING
- cluster_enum: STRING
- datacenter: STRING
- datacenter_enum: STRING
- description: STRING
- description_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- network_name: STRING
- network_name_enum: STRING

- node: STRING
- nodeid: STRING
- physical_addr: STRING
- physical_addr_enum: STRING
- pkts_recd: INTEGER
- pkts_recd_enum: STRING
- pkts_trans: INTEGER
- pkts_trans_enum: STRING
- received: INTEGER
- received_enum: STRING
- switch: STRING
- switch_enum: STRING
- timestamp: STRING
- transmitted: INTEGER
- transmitted_enum: STRING
- vm_hostname: STRING
- vm_hostname_enum: STRING
- vm_name: STRING
- vm_name_enum: STRING
- vm_os_type: INTEGER
- vm_os_type_enum: STRING
- vm_server_name: STRING
- vm_server_name_enum: STRING

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

- datacenter: STRING
- datacenter_enum: STRING
- datastore: STRING
- datastore_cluster: STRING
- datastore_cluster_enum: STRING
- datastore_enum: STRING
- file_path: STRING
- file_path_enum: STRING
- file_size: INTEGER
- file_size_enum: STRING
- kvm_source: STRING
- kvm_source_enum: STRING
- last_modified: STRING
- last_modified_enum: STRING

- node: STRING
- owner: STRING
- owner_enum: STRING
- timestamp: STRING

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

- capacity: INTEGER
- capacity_enum: STRING
- description: STRING
- description_enum: STRING
- free_space: INTEGER
- free_space_enum: STRING
- include_data_in_summarization_0: INTEGER
- include_data_in_summarization_0_enum: STRING
- node: STRING
- nodeid: STRING
- percent_free: INTEGER
- percent_free_enum: STRING
- percent_used: INTEGER
- percent_used_enum: STRING
- timestamp: STRING
- used_space: INTEGER
- used_space_enum: STRING
- vm_hostname: STRING
- vm_hostname_enum: STRING
- vm_name: STRING
- vm_name_enum: STRING
- vm_os_type: INTEGER
- vm_os_type_enum: STRING
- vm_server_name: STRING
- vm_server_name_enum: STRING

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

- node: STRING
- timestamp: STRING

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

- node: STRING
- timestamp: STRING

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

- creation_time: STRING
- creation_timestamp: STRING
- creation_time_enum: STRING
- description: STRING
- description_enum: STRING
- node: STRING
- snapshot_moref: STRING
- snapshot_moref_enum: STRING
- snapshot_name: STRING
- snapshot_name_enum: STRING
- space_consumed: INTEGER
- space_consumed_enum: STRING
- timestamp: STRING
- vm_name: STRING
- vm_name_enum: STRING
- vm_state: STRING
- vm_state_enum: STRING

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

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

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

Trademarks

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

Intel, Intel logo, and Intel Xeon, are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

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

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

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

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's user name for purposes of session management, authentication, and single sign-on configuration. These cookies cannot be disabled.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Index

A

B

C

D

E

F

G

H

I

K

L

M

N

O

P

Q

R

S

T

U

V

W