Monitoring VMware vCenter

eG Innovations Product Documentation

www.eginnovations.com

Total Performance Visibility

Table of Contents

CHAPTER 1: INTRODUCTION 1
CHAPTER 2: HOW DOES EG ENTERPRISE MONITOR VMWARE VCENTER? 3
2.1 Managing the VMware vCenter 3
CHAPTER 3: MONITORING VMWARE VCENTER 5
3.1 The vCenter Services Layer 7
3.1.1 ESX Server Status in vCenter Test 7
3.1.2 Hardware Sensors Test 15
3.1.3 vCenter Status Test 18
3.1.4 vCenter Sessions Test 20
3.1.5 vCenter Licenses Test 24
3.1.6 ESX Servers Test 38
3.1.7 vCenter User Sessions Test 44
3.1.8 vCenter Tasks Test 46
3.1.9 vCenter Events Test 54
3.1.10 vCenter VMotions Test 62
3.1.11 vSphere vRAM Usage Test 65
3.1.12 VM Movement Test 69
3.1.13 vSAN Host Health Status Test 74
3.2 The Virtual Networking Layer 77
3.2.1 Distributed vSwitch Ports Test 77
3.3 The VSAN Layer 83
3.3.1 vSAN Disk Group Test 84
3.3.2 vSAN Health Test 96
3.3.3 vSAN Overview Test 100
3.3.4 vSAN Physical DiskTest 104
3.3.5 vSAN VMKernel Test 110
3.4 The Virtual Machine Clusters Layer 112
3.4.1 Virtual Clusters Test 112
3.4.2 Cluster Resource Pools Test 125
3.5 The Datacenters Layer 137
3.5.1 Datastores Test 137
3.5.2 Datacenters Test 148
CHAPTER 4: MONITORING VCENTER CLUSTERS 152
ABOUT EG INNOVATIONS 153

Table of Figures

Figure 1.1: The architecture of vCenter 1
Figure 2.1: Adding a vCenter server 3
Figure 2.2: List of Unconfigured tests for the vCenter server 4
Figure 3.1: The layer model of vCenter 5
Figure 3.2: The tests mapped to the vCenter Services layer 7
Figure 3.3: The detailed diagnosis of the Warning sensors measure 18
Figure 3.4: The detailed diagnosis of the Sessions logging out measure 24
Figure 3.5: The detailed diagnosis of the Current sessions measure 24
Figure 3.6: Connecting to vCenter 29
Figure 3.7: The VMware Infrastructure Client console 29
Figure 3.8: Adding a new role 30
Figure 3.9: Providing the details of the new role 31
Figure 3.10: The newly created role appended to the list of existing roles in the Roles tab page 32
Figure 3.11: The Permissions tab page 33
Figure 3.12: Selecting the Add Permission option 34
Figure 3.13: Adding a user 35
Figure 3.14: Selecting a local/domain user 36
Figure 3.15: Assign the newly created role to the user 37
Figure 3.16: The Permissions tab page displaying the chosen user and the role assigned to him/her 38
Figure 3.17: The detailed diagnosis of the ESX servers managed by VC measure 43
Figure 3.18: The detailed diagnosis of the ESX servers connected to VC measure 44
Figure 3.19: The detailed diagnosis of the Current sessions to VC measure 46
Figure 3.20: Configuring the tasks to be monitored 47
Figure 3.21: The detailed diagnosis of the successful tasks measure 54
Figure 3.22: Configuring the events to be monitored 55
Figure 3.23: The detailed diagnosis of the Information measure 62
Figure 3.24: The detailed diagnosis of the Error measure 62
Figure 3.25: VMware virtual SAN 74
Figure 3.26: The test mapped to the Virtual Networking layer 77
Figure 3.27: The tests mapped to the VSAN layer 84
Figure 3.28: The detailed diagnosis of the No of datastores measure 103
Figure 3.29: The detailed diagnosis of the No of clusters measure 103
Figure 3.30: The detailed diagnosis of the No of hosts measure 103
Figure 3.31: The tests mapped to the vCenter Cluster layer 112
Figure 3.32: Figure 27: The detailed diagnosis of the Total physical hosts in cluster measure 124
Figure 3.33: The detailed diagnosis of the VMs in cluster measure 124
Figure 3.34: The detailed diagnosis of the VMs powered on measure 124

Figure 3.35: The tests mapped to the Datacenters layer 137
Figure 3.36: The detailed diagnosis of the ESX servers using the datastore measure 147
Figure 3.37: The detailed diagnosis of the VMs using the datastore measure 147
Figure 3.38: The detailed diagnosis of the Esx servers under cluster 151
Figure 3.39: The detailed diagnosis of the Directly connected vms measure 151
Figure 3.40: The detailed diagnosis of the VMs under cluster measure 151
Figure 4.1: The VMware vCenter Cluster monitoring model 152

Chapter 1: Introduction

VMware vCenter delivers centralized management, resource optimization, operational automation and security to virtualized IT environments. The tool serves as a central point of control, management, and configuration of the ESX servers in an environment, and can create, clone, or copy virtual machines/virtual machine templates on the physical servers. Besides being a VM creator, vCenter also dons the role of a resource allocator that allocates memory and CPU resources to virtual machines. By continuously monitoring resource utilization across resource pools, vCenter recognizes a VM's need for resources and dynamically allocates resources to it based on pre-defined rules. Moreover, vCenter can also trigger *VMotion* activity when in need, and can perform live VM migration from one physical server to another.

Figure 1.1 depicts the architecture of vCenter.

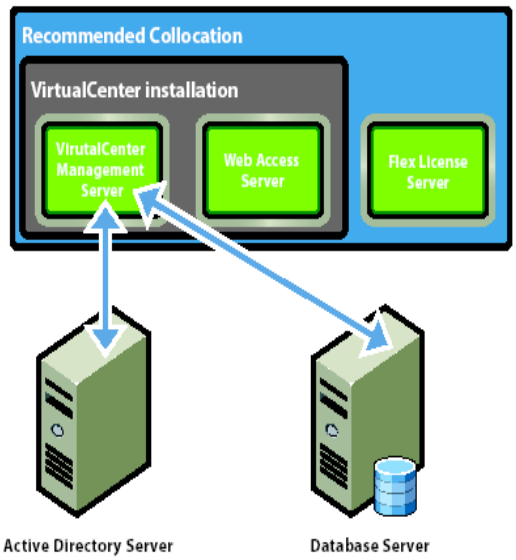

Figure 1.1: The architecture of vCenter

The **vCenter Management Server** is the central control node for configuring, provisioning and managing virtualized IT environments. **vCenter Database** is used to store persistent information about the physical servers, resource pools and virtual machines managed by the vCenter Management Server. **Virtual Infrastructure Web access** allows virtual machine management and access to virtual machine graphical consoles without installing a client. vCenter manages all VMware software licenses with an embedded **FlexNet licensing server** and a single license file. While **Virtual Infrastructure Client** allows administrators and users to connect remotely to the vCenter, the **Active Directory server** is used for providing network-based authentication to these

requests. The **vCenter Agent** connects VMware ESX Server with the vCenter Management Server.

vCenter relies heavily on each of these components, and the failure of any one of these components can either bring vCenter operations to a halt, or force the software to offer only limited functionality. In a "virtual" world, where critical end-user services ride on virtualized components, the non-availability of vCenter, even for a brief while, could prove to be fatal! For that matter, even relatively minor issues, such as insufficient VMware ESX server licenses or excessive resource usage by one/more resource pools, can cause irreparable damage to service delivery, if not continuously tracked and promptly reported. It is therefore imperative that vCenter be monitored 24x7. This is where eG Enterprise helps administrators.

Chapter 2: How does eG Enterprise Monitor VMware vCenter?

eG Enterprise is capable of monitoring the VMware vCenter in both agent based and agentless manners. Before staring to monitor the VMware vCenter, make sure that the eG remote agent uses JRE 1.7. By default, Linux agents are bundled with JRE 1.6. When using such an agent to monitor a VMware infrastructure, ensure that the JRE of the eG agent is upgraded to v1.7.

Note:

VMware vCenter can be monitored only using a 64-bit eG agent.

2.1 Managing the VMware vCenter

1. Log into the eG administrative interface.
2. eG Enterprise cannot automatically discover the VMware vCenter. You need to manually add vCenter using the **COMPONENTS** page (see Figure 2.1) that appears when the Infrastructure -> Components -> Add/Modify menu sequence is followed. Remember that components manually added are managed automatically.

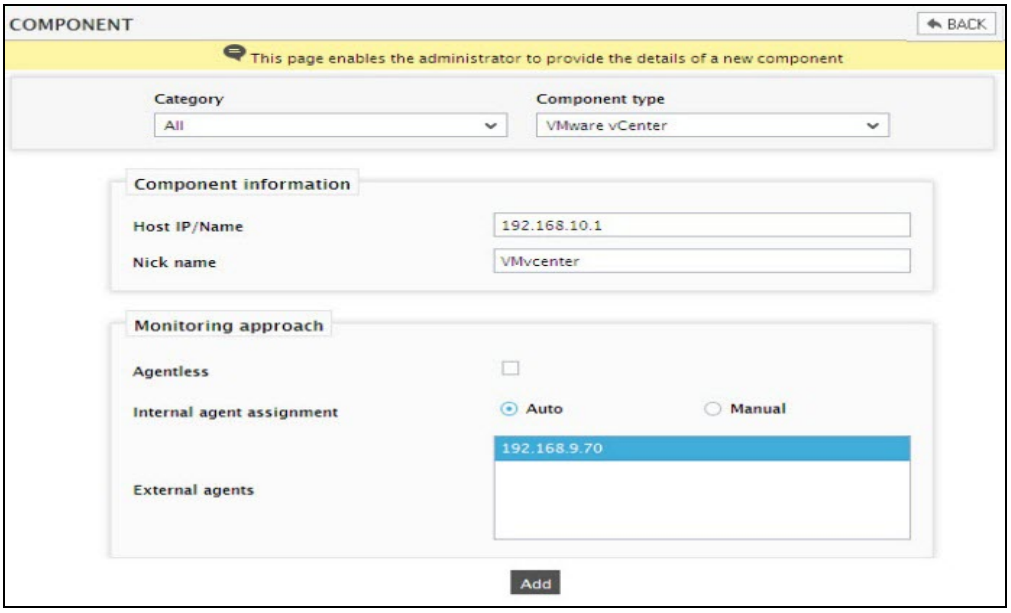

Figure 2.1: Adding a vCenter server

3. When you attempt to sign out, a list of unconfigured tests appears (see Figure 2.2).

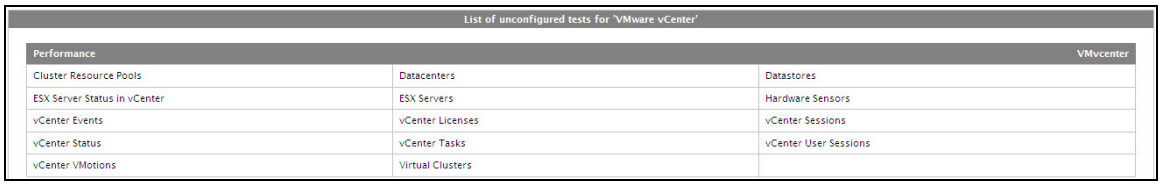

List of unconfigured tests for 'VMware vCenter'

Performance		VMvcenter
Cluster Resource Pools	Datacenters	Datastores
ESX Server Status in vCenter	ESX Servers	Hardware Sensors
vCenter Events	vCenter Licenses	vCenter Sessions
vCenter Status	vCenter Tasks	vCenter User Sessions
vCenter VMotions	Virtual Clusters	

Figure 2.2: List of Unconfigured tests for the vCenter server

4. Click on any test from the list of unconfigured tests to configure it. To know how to configure the parameters for the tests, refer to **Monitoring VMware vCenter**.
5. Finally, sign out of the eG administrative interface.

Chapter 3: Monitoring VMware vCenter

eG Enterprise prescribes a specialized VMware vCenter monitoring model (see Figure 3.1), which can be managed in an agent-based or an agentless manner, and can be configured to periodically check the health of the critical services offered by vCenter, so as to proactively alert administrators to real/potential anomalies.

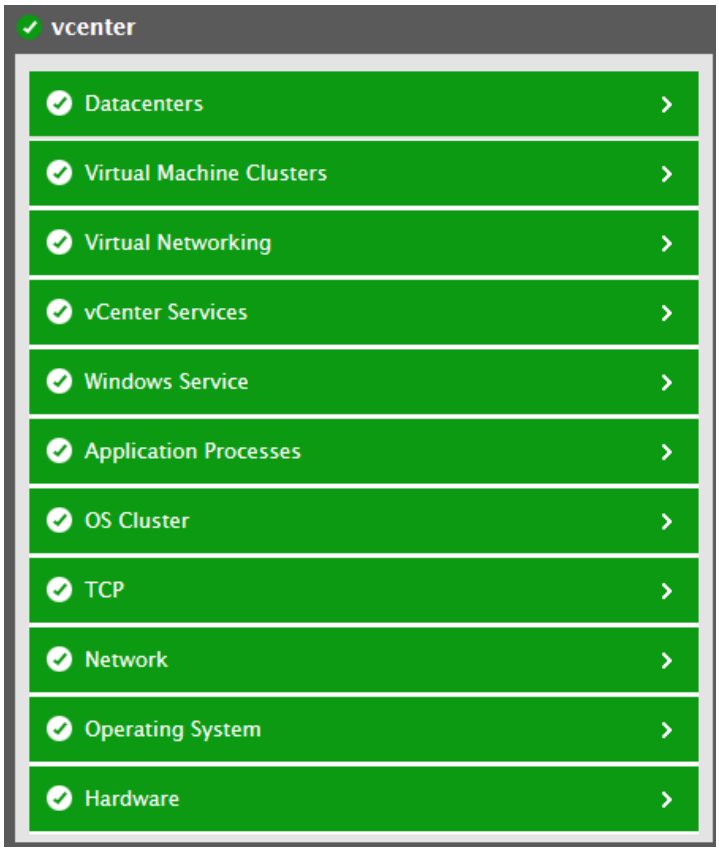

Figure 3.1: The layer model of vCenter

Every layer depicted by Figure 2 is mapped to a wide variety of tests that report the following:

- Is vCenter available? If so, how quickly is it responding to requests?
- Does vCenter possess adequate ESX server licenses?
- Is vCenter overloaded with sessions?
- Are sessions logging out smoothly, or are there any unexpected log outs? Which are the user sessions that logged out suddenly?

- How may sessions does a user have open on VC? Which user has the maximum sessions open on VC? Can I know the details of these sessions?
- Have any VC tasks failed? Which ones are they?
- Which are the tasks that have succeeded?
- Is any VC task taking too long to complete?
- Are there any pending tasks on VC?
- How many ESX servers is VC currently managing? How many of these were added recently?
- How many ESX servers were removed from VC?
- Are any ESX servers disconnected from VC? Were there any recent disconnects?
- Are any ESX servers not responding to VC? Did any ESX server become unresponsive recently?
- How many clusters have been configured on VC? How many ESX servers are managed per cluster?
- How many VMs are available in each cluster, and what is their current state? Are there any powered-off/suspended VMs in a cluster? If so, which ones are they?
- How are the clusters using the resources allocated to them?
- Is any cluster consuming resources excessively? Which resource pool on the cluster is responsible for this? How many VMs and child resource pools does this resource pool consist of?
- What are the datacenters managed by VC? What is the current configuration of each datacenter – i.e., how many ESX servers exist with a datacenter? How many VMs, clusters, networks, and datastores are available for the datacenter?
- How many ESX servers and VMs are directly assigned to a datacenter, and how many operate within a cluster?
- Which datastores have been configured for use of the virtual infrastructure? Which ones are currently available?
- How many ESX servers and VMs are using a particular datastore? Which ones are they?
- Is adequate free space available with the datastore?

The answers to the above questions can shed light on many key performance aspects of vCenter, and enable administrators to accurately pin-point the root-cause of problems with vCenter.

The sections to come discuss each of the layers in Figure 2 elaborately. As the bottom 7 layers of the monitoring model have already been dealt with in the *Monitoring Unix and Windows Servers* document, the sections to come will deal with the remaining layers alone.

3.1 The vCenter Services Layer

The tests mapped to this layer perform the following functions:

- Report the availability and responsiveness of vCenter
- Monitor the sessions to vCenter, including sessions per user
- Monitor how vCenter manages licenses
- Track the status of tasks and events executing on VC
- Report the status of ESX hosts managed by VC
- Track and report the status of VMotion activity

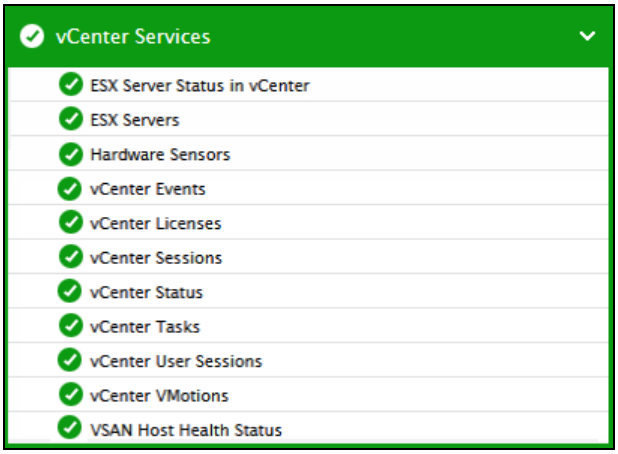

Figure 3.2: The tests mapped to the vCenter Services layer

3.1.1 ESX Server Status in vCenter Test

This test reports the current state of each ESX server managed by vCenter. Disconnected hosts, hosts in the maintenance mode, and powered off hosts can be quickly identified using this test.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each vSphere/ESX host managed by the vCenter server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is

Parameter	Description
	set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Connection status	Indicates the current connection state of this vSphere/ESX server.		The values this measure reports and their numeric equivalents are provided in the table below: **Measure Value** / **Numeric Value** Disconnected: 0 Connected: 1 NotResponding: 2 **Note:** Typically, this measure reports one of the **Measure Value**s listed in the table above. In the graph of this measure however, the connection state of an ESX server is indicated by its corresponding numeric equivalent only - i.e., 0 to 2.
State of an active host in HA cluster	Indicates the current availability state of this host in an HA cluster.		In a vSphere HA cluster, the active hosts form a fault domain. A host is inactive if it is in standby or maintenance mode, or it has been disconnected from vCenter Server. A vSphere HA agent, called the Fault Domain Manager (FDM), runs on each host in the fault domain. One FDM serves as the master and the remaining FDMs as its slaves. The master is responsible for monitoring the availability of the hosts and VMs in the

Measurement	Description	Measurement Unit	Interpretation
			cluster, and restarting any VMs that fail due to a host failure or non-user-initiated power offs. The master is also responsible for reporting fault-domain state to vCenter Server. The master FDM is determined through election by the FDMs that are alive at the time. An election occurs in the following circumstances: a. When the vSphere HA feature is enabled for the cluster. b. When the master's host fails. c. When the management network is partitioned. In a network partition there will be a master for each partition. However, only one master will be responsible for a given VM. When the partition is resolved, all but one of the masters will abdicate. d. After a host in a vSphere HA cluster powers back up following a failure that caused all hosts in the cluster to power off. The slaves are responsible for reporting state updates to the master and restarting VMs as required. The values that this measure reports and their numeric equivalents have been discussed in the table below

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Master</td><td>0</td></tr>
<tr><td>ConnectedToMaster</td><td>1</td></tr>
<tr><td>Election</td><td>2</td></tr>
<tr><td>Uninitialized</td><td>3</td></tr>
<tr><td>NetworkPartitioned FromMaster</td><td>4</td></tr>
<tr><td>FdmUnreachable</td><td>5</td></tr>
<tr><td>HostDown</td><td>6</td></tr>
<tr><td>NetworkIsolated</td><td>7</td></tr>
<tr><td>InitializationError</td><td>8</td></tr>
<tr><td>UnitializationError</td><td>9</td></tr>
</table>
Note:
Typically, this measure reports one of the Measure Values listed in the table above. In the graph of this measure however, the availability state of an ESX server in an HA cluster is indicated by its corresponding numeric equivalent only - i.e., 0 to 9.
</td></tr>
<tr><td>Is in maintenance mode?</td><td>Indicates whether this host is currently in the maintenance mode or not.</td><td></td><td>If a vSphere host is in the maintenance mode, then VMware HA suspends failover operations during maintenance operations. In other words, when a host enters the maintenance mode, VMs are prevented from powering up or failing over to the host if that host is taking part in a high availability cluster.
The values this measure reports and their numeric equivalents are provided in the table below:</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Yes</td><td>1</td></tr>
<tr><td>No</td><td>0</td></tr>
</table>
Note:
Typically, this measure reports one of the Measure Values listed in the table above. In the graph of this measure however, the maintenance mode of an ESX server is indicated by its corresponding numeric equivalent only - i.e., 0 or 1.
</td></tr>
<tr><td>Standby mode</td><td>Indicates whether this host is currently in the standby mode or not.</td><td></td><td>The VMware Distributed Power Management (DPM) feature allows a DRS cluster to reduce its power consumption by powering hosts on and off based on cluster resource utilization. VMware DPM monitors the cumulative demand of all virtual machines in the cluster for memory and CPU resources and compares this to the total available resource capacity of all hosts in the cluster. If sufficient excess capacity is found, VMware DPM migrates all the virtual machines on one/more hosts to other hosts and places such hosts in the standby mode (i.e., powers off the hosts). Conversely, when capacity is deemed to be inadequate, DRS brings hosts out of standby mode (powers them on) and migrates virtual machines, using VMotion, to them.
The values this measure reports and their numeric equivalents are provided in the table below:</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>None</td><td>0</td></tr>
<tr><td>In</td><td>1</td></tr>
<tr><td>Entering</td><td>2</td></tr>
<tr><td>Exiting</td><td>3</td></tr>
</table>
Note:
Typically, this measure reports one of the Measure Values listed in the table above. In the graph of this measure however, the standby mode of an ESX server is indicated by its corresponding numeric equivalent only - i.e., 0 to 3.
</td></tr>
<tr><td>Power state</td><td>Indicates whether this host is currently in the standby, powered off, or powered on mode.</td><td></td><td>The values this measure reports and their numeric equivalents are provided in the table below:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Unknown</td><td>0</td></tr>
<tr><td>PoweredOn</td><td>1</td></tr>
<tr><td>PoweredOff</td><td>2</td></tr>
<tr><td>StandBy</td><td>3</td></tr>
</table>
Note:
Typically, this measure reports one of the Measure Values listed in the table above. In the graph of this measure however, the power state of an ESX server is indicated by its corresponding numeric equivalent only - i.e., 0 to 3.
</td></tr>
<tr><td>Availability status</td><td>Indicates whether this host is currently in the maintenance, standby, powered on, or powered off modes.</td><td></td><td>The values this measure reports and their numeric equivalents are provided in the table below:</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>In maintenance</td><td>0</td></tr>
<tr><td>PoweredOn</td><td>1</td></tr>
<tr><td>Unknown</td><td>2</td></tr>
<tr><td>PoweredOff</td><td>3</td></tr>
<tr><td>StandBy</td><td>4</td></tr>
</table>
Note:
Typically, this measure reports one of the Measure Values listed in the table above. In the graph of this measure however, the availability state of an ESX server is indicated by its corresponding numeric equivalent only - i.e., 0 to 4.
</td></tr>
<tr><td>Is vSAN enabled?</td><td>Indicates whether/not the vSAN service is currently enabled for this host.</td><td></td><td>Virtual SAN virtualizes local physical storage resources of vSphere hosts and turns them into pools of storage that can be carved up and assigned to virtual machines and applications according to their quality of service requirements.
You can activate Virtual SAN when you create host clusters or enable Virtual SAN on existing clusters. When enabled, Virtual SAN aggregates all local storage disks available on the hosts into a single datastore shared by all hosts. You can later expand the datastore by adding storage devices or hosts to the cluster.
If a host contributes its local storage to the Virtual SAN datastore, the host must provide one SSD and at least one HDD, also called data disk. The disks on the contributing host form a disk group. Each disk group must include one SSD and at least one or multiple data disks. The disk</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>group uses the SSD disk for read caching and write buffering, while the data disks are used for persistent storage. You can have multiple disk groups per host.

If vSAN is not enabled for a host, then the value of this measure will be No. If vSAN is enabled for a host, then the value of this measure will be Yes. The numeric values that correspond to these measure values are discussed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

By default, this measure reports the Measure Values listed in the table above. In the graph of this measure however, the state of the vSAN is represented using the corresponding numeric equivalents only.</td></tr>
</table>

3.1.2 Hardware Sensors Test

This test reports the number and nature of the hardware sensors operating on each vSphere/ESX host that is managed by the vCenter server.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each vSphere/ESX host managed by the vCenter server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is

Parameter	Description
	set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Descrption	Measurement Unit	Interpretation
Total sensors	Indicates the total number of sensors available in this vSphere/ESX host.	Number	
Warning sensors	Indicates the total number of sensors in this vSphere/ESX host that are currently in the warning state.	Number	Use the detailed diagnosis of this measure to know which sensors are in the warning state.
Percentage of warning sensors	Indicates the percentage of sensors in this vSphere/ESX host that are currently in a warning state.	Percent	Ideally, the value of this measure should be 0. A non-zero value indicates the probable failure of one/more sensor.
Error sensors	Indicates the number of sensors in this	Number	Use the detailed diagnosis of this measure to know which sensors are

Measurement	Descrption	Measurement Unit	Interpretation
	vSphere/ESX host with errors.		currently experiencing errors.
Percentage of error sensors	Indicates the percentage of sensors in this vSphere/ESX host that are currently experiencing errors.	Percent	Ideally, the value of this measure should be 0. A non-zero value indicates that one/more sensors have encountered errors while operating.

Use the detailed diagnosis of the *Warning sensors* measure to know which sensors are in the warning state.

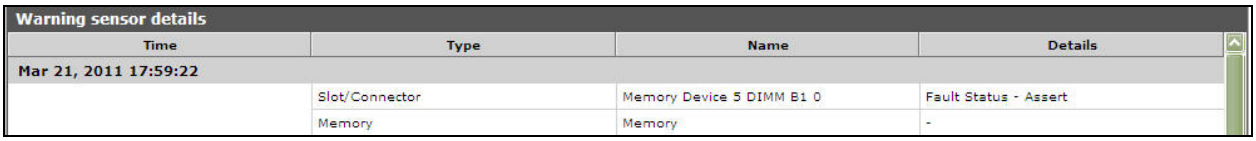

Figure 3.3: The detailed diagnosis of the Warning sensors measure

3.1.3 vCenter Status Test

This test monitors the availability and response time of vCenter.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the vCenter server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
Using	The Using parameter determines whether this test should be monitored using the web services API or CLI (command line interface). If you are monitoring vCenter in an agentless manner, then you might want to execute this test using the web services API, as this approach supports both agent-based and

Parameter	Description
	agentless monitoring modes. In such a case therefore, select **Web services API** from the Using list box. If the **Web services API** option is chosen, then ensure that the Virtual Center Server Home parameter is set to *none*. Also, make sure that the VC User and VC Password parameters represent the user name and password to connect to the vCenter server. By default, **Web Services API** is selected from the Using list. The command line interface is typically preferred, if vCenter is monitored in an agent-based manner. In such a case therefore, select **Command Line Interface** from the Using list. If this option is chosen, then ensure that the VirtualCenter License File, Virtual Center Server Home and ReReadLicense parameters are configured with valid values, and the VC User and VC Password parameters are set to *None*. Note: vCenter 4.0 does not support the command line interface; therefore, set the **USING** parameter to **Web services API** while monitoring vCenter 4.0.
VC User and VC Password	If this test uses the **Web Services API**, then in the VC User and VC Password text boxes, specify the name and password (respectively) of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if owing to security constraints you prefer not to expose the credentials of such a user, then, you can create a special role for this purpose on vCenter, assign the **Licenses** permission to this role, and assign this role to a local/domain user to vCenter. The procedure for creating this special role and assigning it to an existing local/domain user has been detailed in Section 1.1.3.1 of this document. On other hand, if this test executes using the **Command Line Interface**, then the VC user and vc password parameters should be set to *none*.
Confirm Password	Confirm the password by retyping it in this text box.
Virtual Center Server Home	If the using parameter is set to **Web services API**, then specify *none* against Virtual Center Server Home. On the other hand, if **Command Line Interface** is chosen from the Using list, then provide the full path to the install directory of the VMware Licensing server. For example, *c:\progra~1\VMware\Vmware License Server*.
VirtualCenter License File	In the VirtualCenter License File text box, specify the name of the vCenter license file used by this test for reporting license-related statistics; this is *vmware.lic* by default. This parameter is of relevance only if the Using flag is set to **Command Line Interface**.
ReReadLicense	If the test executes using the **Web services API**, then you will have to set this flag to true or false to indicate whether the eG agent should or should not check for license

Parameter	Description
	changes everytime it runs this test. If this flag is set to **true**, then the eG agent will check for changes in license status everytime this test runs. If this flag is set to **false**, then the eG agent will not check for license changes. The status of the ReReadLicense flag will be disregarded if the using parameter is set to **Command Line Interface**.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the SSL flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Availability	Indicates whether the vCenter server is available or not.	Percent	The value 0 indicates that the vCenter is not available; whereas, the value 100 indicates that it is available.
Response time	Indicates the time taken by the vCenter to respond to requests.	Secs	If the value of this measure consistently increases, it indicates that the performance of the vCenter is gradually deteriorating. Ideally, the value of this measure should be low.

3.1.4 vCenter Sessions Test

This test monitors the user sessions to the vCenter, and reports new logins and sessions logging out.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the vCenter server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
Using	The Using parameter determines whether this test should be monitored using the web services API or CLI (command line interface). If you are monitoring vCenter in an agentless manner, then you might want to execute this test using the web services API, as this approach supports both agent-based and agentless monitoring modes. In such a case therefore, select **Web services API** from the Using list box. If the **Web services API** option is chosen, then ensure that the Virtual Center Server Home parameter is set to *none*. Also, make sure that the VC User and VC Password parameters represent the user name and password to connect to the vCenter server. By default, **Web Services API** is selected from the Using list. The command line interface is typically preferred, if vCenter is monitored in an agent-based manner. In such a case therefore, select **Command Line Interface** from the Using list. If this option is chosen, then ensure that the VirtualCenter License File, Virtual Center Server Home and ReReadLicense parameters are configured with valid values, and the VC User and VC Password parameters are set to *None*. Note: vCenter 4.0 does not support the command line interface; therefore, set the **USING** parameter to **Web services API** while monitoring vCenter 4.0.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document.

Parameter	Description
	vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the SSL flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability

Parameter	Description
	• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current sessions to VC	Indicates the total number of logins to vCenter.	Number	Use the detailed diagnosis of this measure to know the details of the current sessions.
New logins to VC	Indicates the number of new logins to vCenter.	Number	
Percent new logins	Indicates the percentage of current sessions that logged in during the last measurement period.	Percent	
Sessions logging out	Indicates the number of sessions that logged out.	Number	If all the current sessions suddenly log out, it indicates a problem condition that requires investigation. The detailed diagnosis of this measure, if enabled, reveals the details of the sessions that logged out.

The detailed diagnosis of the *Sessions logging out* measure reveals the users who initiated the sessions, the time of login, and the duration of the sessions.

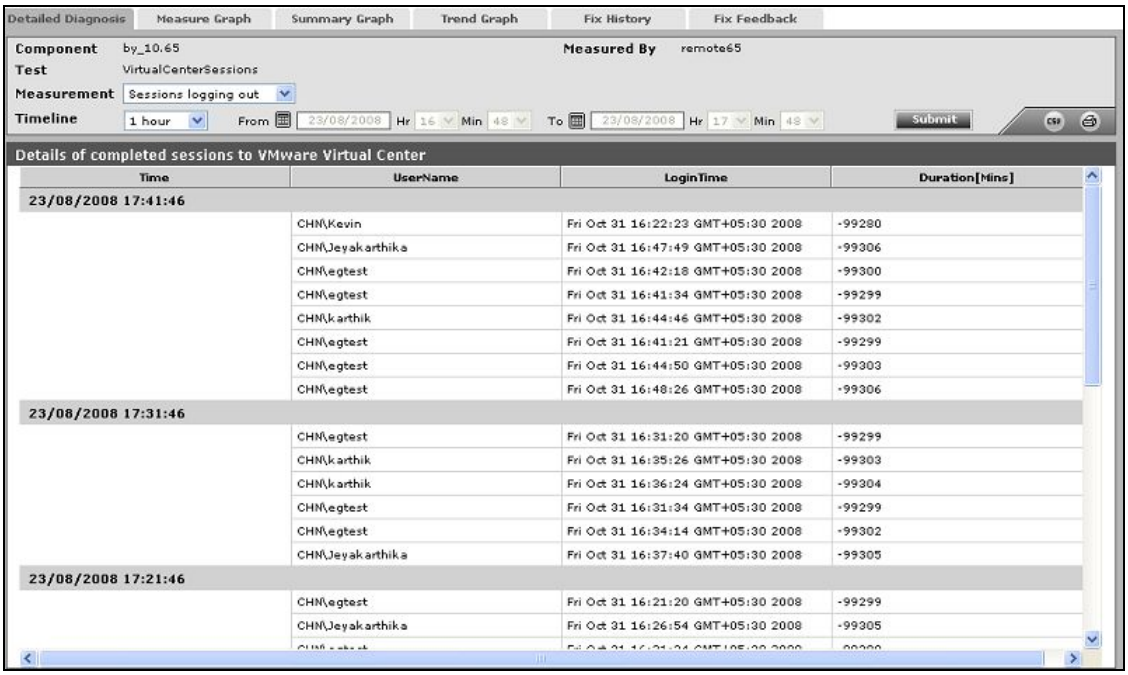

Figure 3.4: The detailed diagnosis of the Sessions logging out measure

The detailed diagnosis of the *Current sessions* measure reveals the details of these sessions.

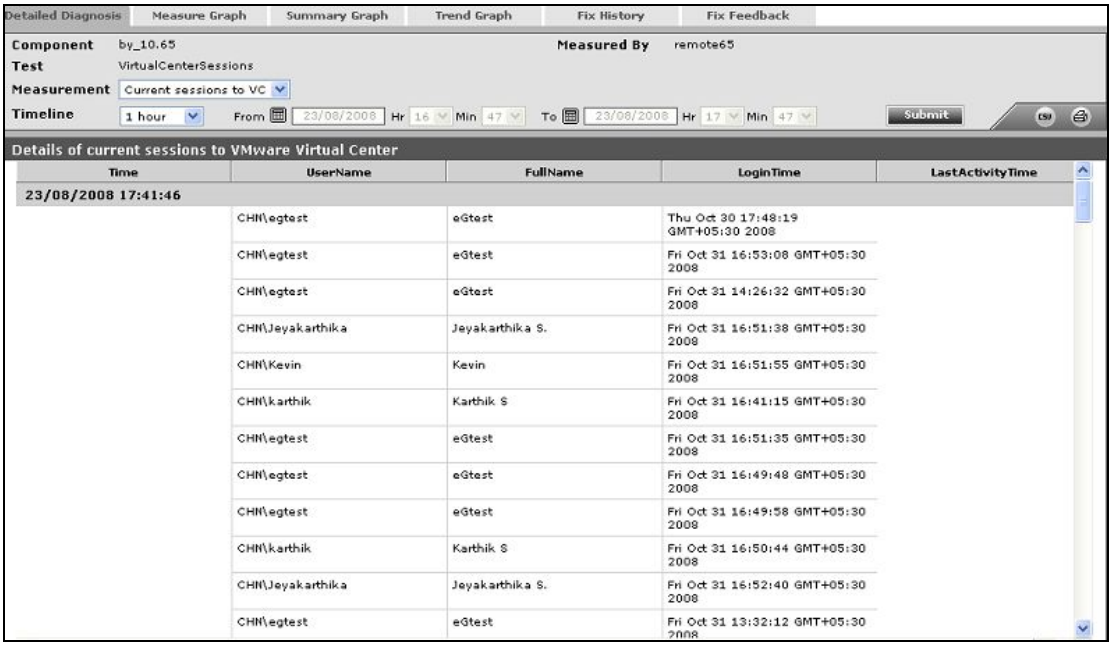

Figure 3.5: The detailed diagnosis of the Current sessions measure

3.1.5 vCenter Licenses Test

This test reports critical statistics related to the license management by vCenter.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the vCenter server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
Using	The Using parameter determines whether this test should be monitored using the web services API or CLI (command line interface). If you are monitoring vCenter in an agentless manner, then you might want to execute this test using the web services API, as this approach supports both agent-based and agentless monitoring modes. In such a case therefore, select **Web services API** from the Using list box. If the **Web services API** option is chosen, then ensure that the Virtual Center Server Home parameter is set to *none*. Also, make sure that the VC User and VC Password parameters represent the user name and password to connect to the vCenter server. By default, **Web Services API** is selected from the Using list. The command line interface is typically preferred, if vCenter is monitored in an agent-based manner. In such a case therefore, select **Command Line Interface** from the Using list. If this option is chosen, then ensure that the VirtualCenter License File, Virtual Center Server Home and ReReadLicense parameters are configured with valid values, and the VC User and VC Password parameters are set to *None*. **Note:** vCenter 4.0 does not support the command line interface; therefore, set the **USING** parameter to **Web services API** while monitoring vCenter 4.0.
VC User and VC Password	If this test uses the **Web Services API**, then in the VC User and VC Password text boxes, specify the name and password (respectively) of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if owing to security constraints you prefer not to expose the credentials of such a user, then, you can create a special role for this purpose on vCenter, assign the **Licenses** permission to this role, and assign this role to a local/domain user to vCenter. The procedure for creating this special role and assigning it to an existing local/domain user has been detailed in Section 1.1.3.1 of this document.

Parameter	Description
	On other hand, if this test executes using the **Command Line Interface**, then the VC user and vc password parameters should be set to *none*.
Confirm Password	Confirm the password by retyping it in this text box.
Virtual Center Server Home	If the using parameter is set to **Web services API**, then specify *none* against Virtual Center Server Home. On the other hand, if **Command Line Interface** is chosen from the Using list, then provide the full path to the install directory of the VMware Licensing server. For example, *c:\progra~1\VMware\Vmware License Server*.
VirtualCenter License File	In the VirtualCenter License File text box, specify the name of the vCenter license file used by this test for reporting license-related statistics; this is *vmware.lic* by default. This parameter is of relevance only if the Using flag is set to **Command Line Interface**.
ReReadLicense	If the test executes using the **Web services API**, then you will have to set this flag to true or false to indicate whether the eG agent should or should not check for license changes everytime it runs this test. If this flag is set to **true**, then the eG agent will check for changes in license status everytime this test runs. If this flag is set to **false**, then the eG agent will not check for license changes. The status of the ReReadLicense flag will be disregarded if the using parameter is set to **Command Line Interface**.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the SSL flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option.

Parameter	Description
	The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Licenses installed	Indicates the number of licenses currently installed on vCenter.	Number	
Licenses in use	Indicates the number of licenses currently being used.	Number	Ideally, this value should be lesser than the value of the *Licenses installed* measure. If this measure is equal to Licenses installed, then it indicates that the vCenter has run out of licenses. Use the detailed diagnosis capability of this measure to view the details of the licenses that are in use. **Note:** If the using parameter of this test has been set to **Web Services API,** then the eG agent will not be able to generate detailed measures for this test – i.e., the detailed diagnosis for this measure will not be available. If the using parameter is set to **Command Line Interface**, then you will be able to view the detailed diagnosis of this measure.
Available licenses	Indicates the number of licenses not in use currently.	Number	
License utilization	Indicates the percentage	Percent	Ideally, this value should be low. A

Measurement	Description	Measurement Unit	Interpretation
	of licenses currently in use.		value equal to or close to 100% is a cause for concern, as it indicates that the vCenter has run out of licenses.

3.1.5.1 Configuring a Special Role on vCenter and Assigning the Role to an Existing Local/Domain User

If the **vCenter Licenses** test has been configured to use the **Web Services API** for collecting license statistics from vCenter, then the eG agent executing the test should possess either **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, owing to security constraints, users of some environments might not want to expose the credentials of their **Administrators** or **Virtual Machine Administrators** . In such environments therefore, administrators can create a special role for this purpose on vCenter, and assign that role to an existing local/domain user. The steps for the same are detailed below:

1. Login to a system on which the **VMware Infrastructure Client** is installed.
2. Double-click on the **VMware Infrastructure Client** icon on your desktop.
3. Figure 3.6 then appears. To connect to the vCenter, select the **IP address / Name** of the vCenter, and then provide the login information. To grant access permissions to a user, you will have to login to vCenter as an *administrator*. Therefore, provide the **User name** and **Password** of the administrator in Figure 3.6.

Figure 3.6: Connecting to vCenter

4. Figure 7 then appears. Click on the **Administration** icon indicated by Figure 3.7.

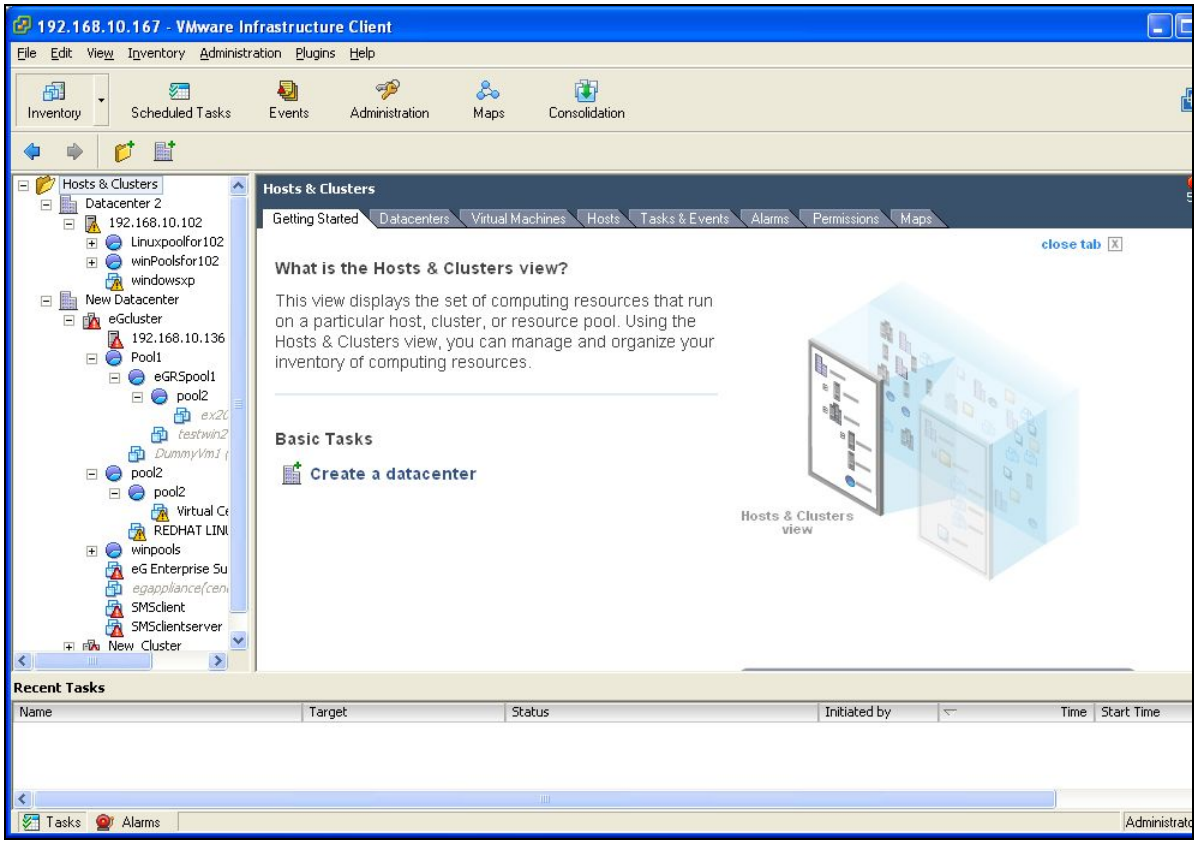

Figure 3.7: The VMware Infrastructure Client console

5. Doing so opens the **Roles** tab page by default. This lists all the existing user roles on vCenter. To create

a new role, click on the **Add Role** button indicated by Figure 3.8.

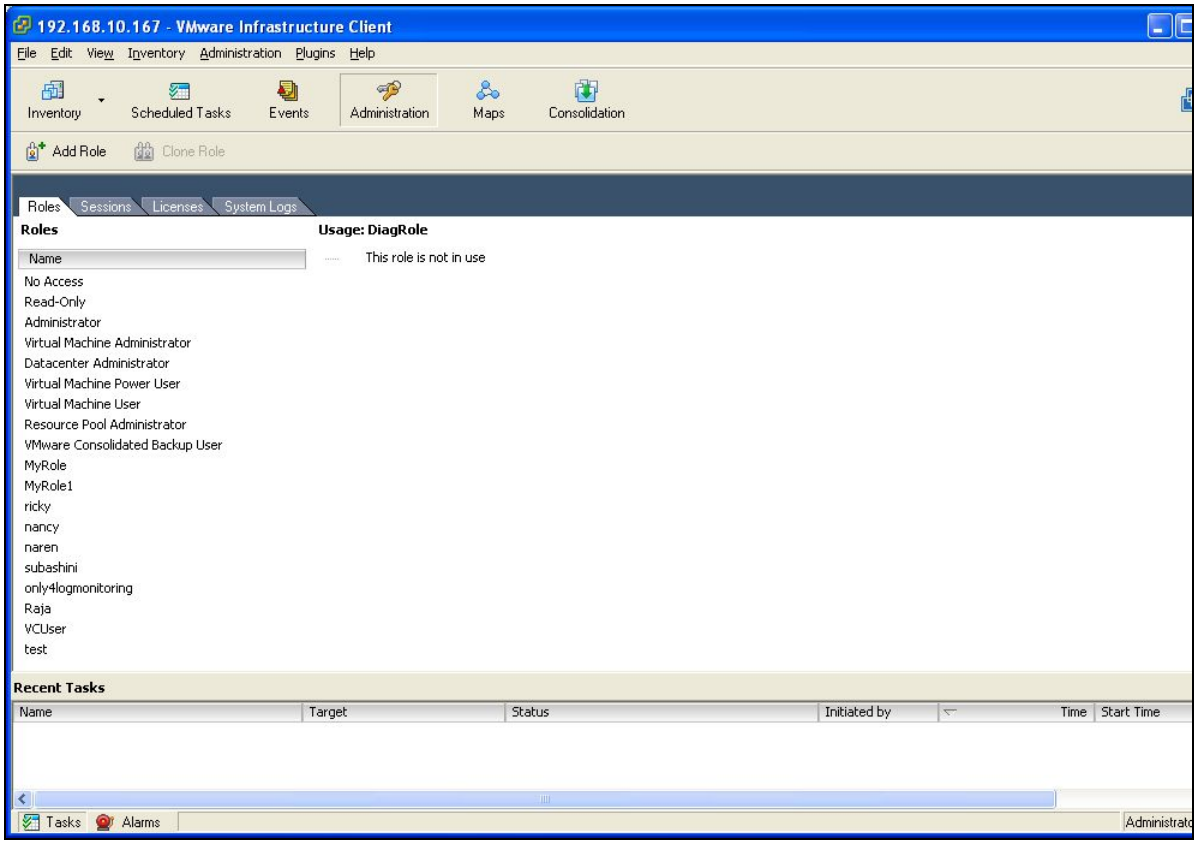

Figure 3.8: Adding a new role

6. In Figure 3.9 that then appears, Enter **Name** of the new role. Then, grant permissions to the new role by first expanding the **Global** node of the tree structure in the **Privileges** section of Figure 3.9. To grant the **Licenses** permission to the new role, select the check box corresponding to **Licenses** in Figure 3.9. Then, click the **OK** button.

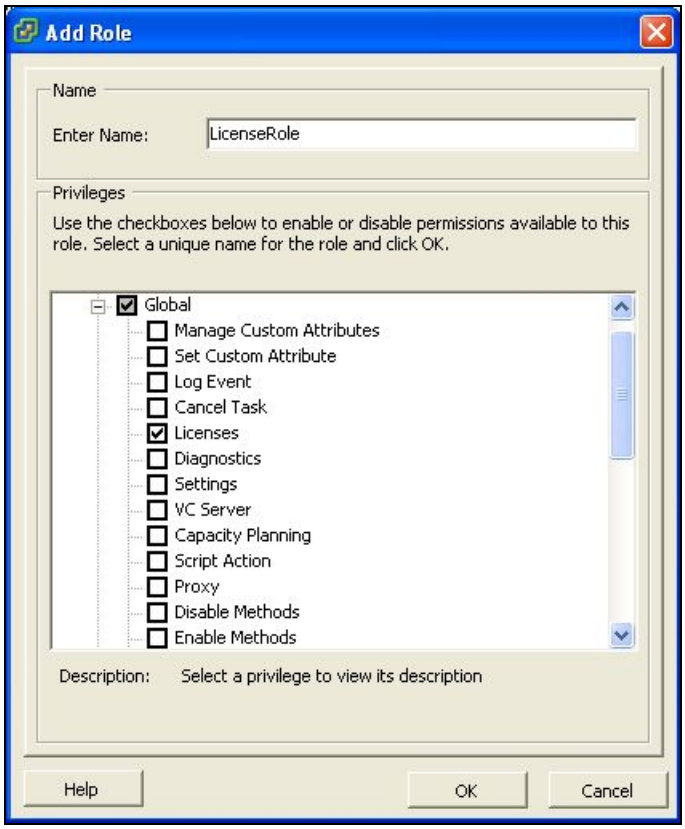

Figure 3.9: Providing the details of the new role

7. Figure 3.10 will then appear displaying the role that you just created.

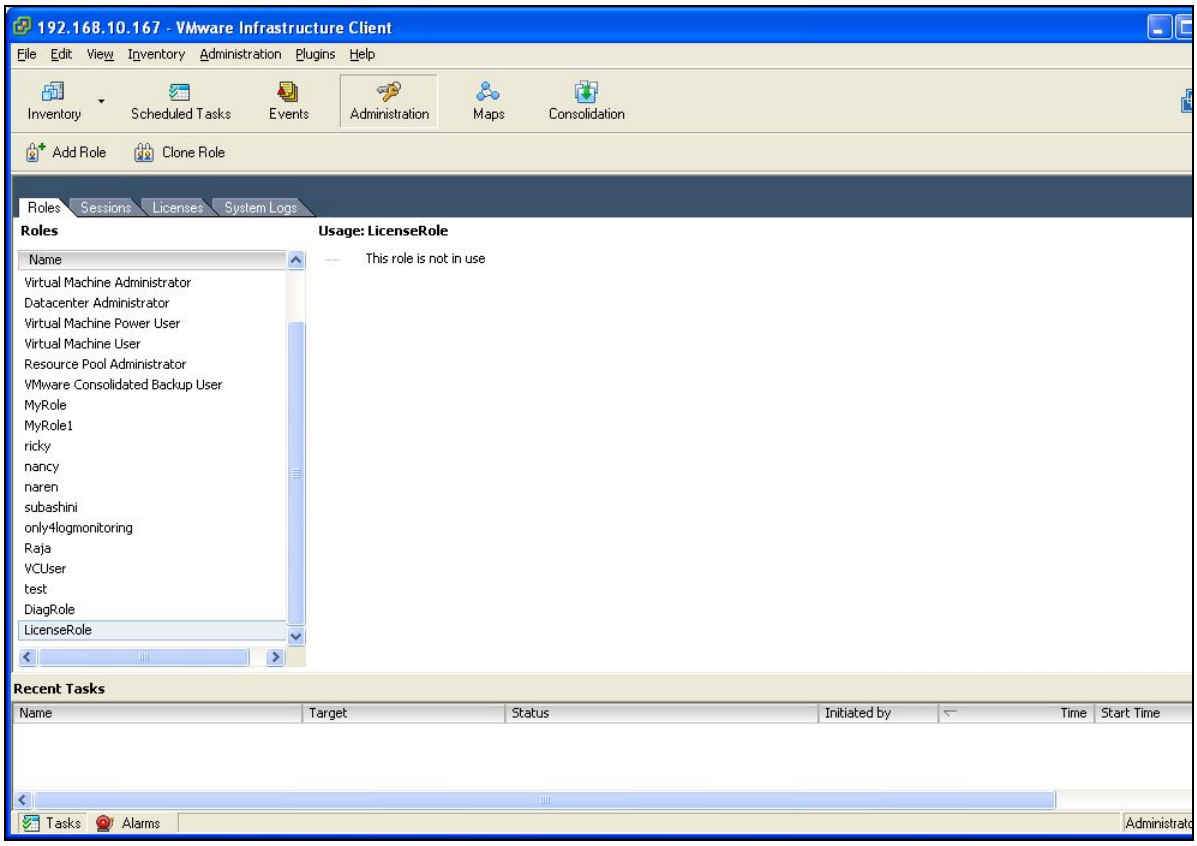

Figure 3.10: The newly created role appended to the list of existing roles in the Roles tab page

8. Next, click on the **Inventory** icon indicated by Figure 3.10. When Figure 3.11 appears, click on the **Hosts & Clusters** node in the tree-structure in the left panel of Figure 3.11, and then click on the **Permissions** tab page in the right panel. This tab page lists the local/domain users to vCenter and the roles currently assigned to them.

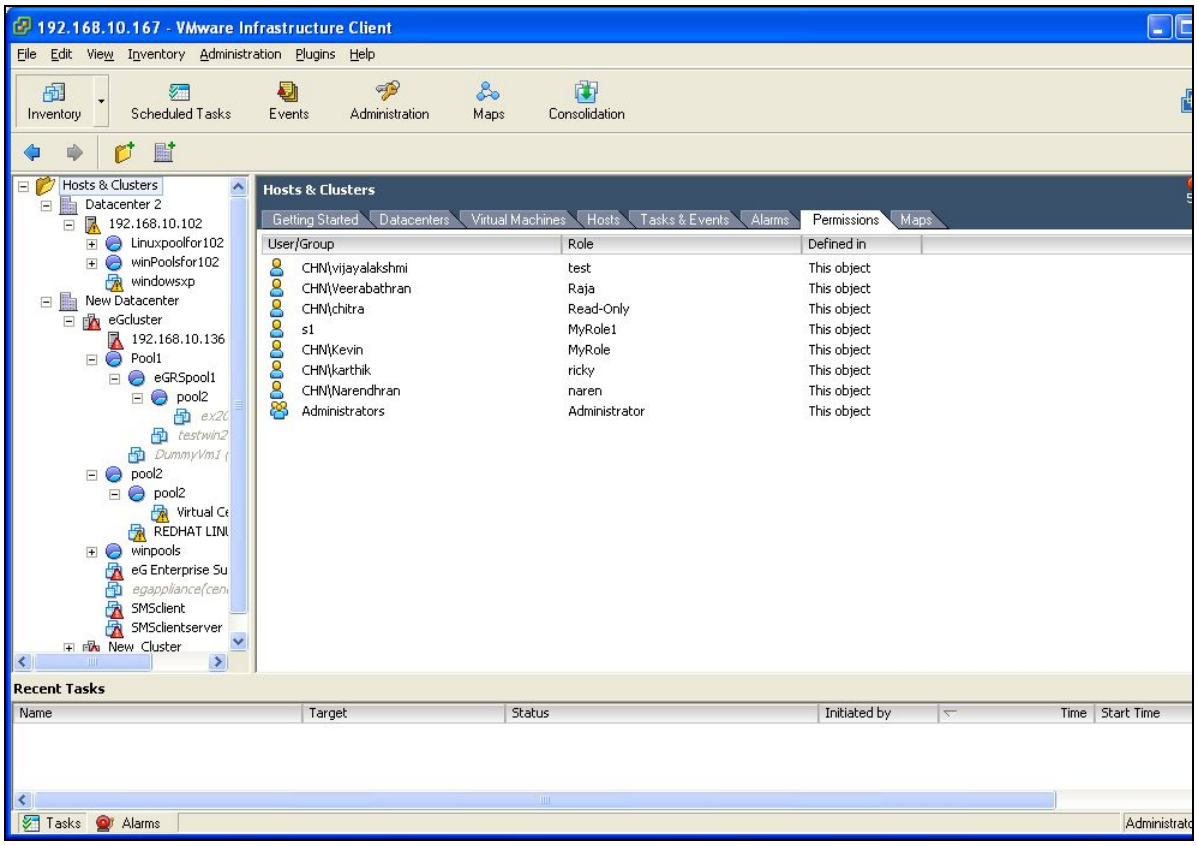

Figure 3.11: The Permissions tab page

9. To assign the newly created role to one of the listed users, right-click anywhere in the empty, white space in the right panel, and select the **Add Permission** option from the shortcut menu that appears.

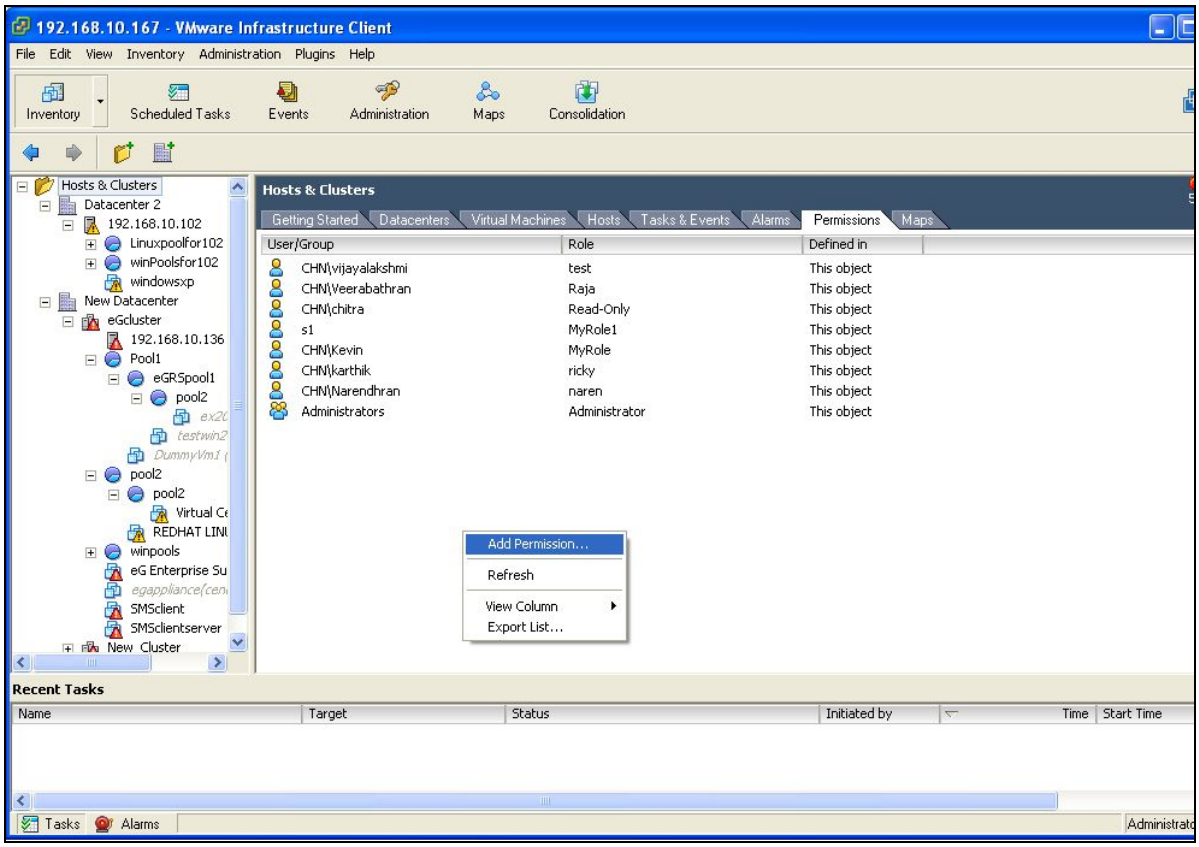

Figure 3.12: Selecting the Add Permission option

10. Figure 3.13 then appears. To choose the user who is to be assigned the new role, click the **Add** button in Figure 3.13.

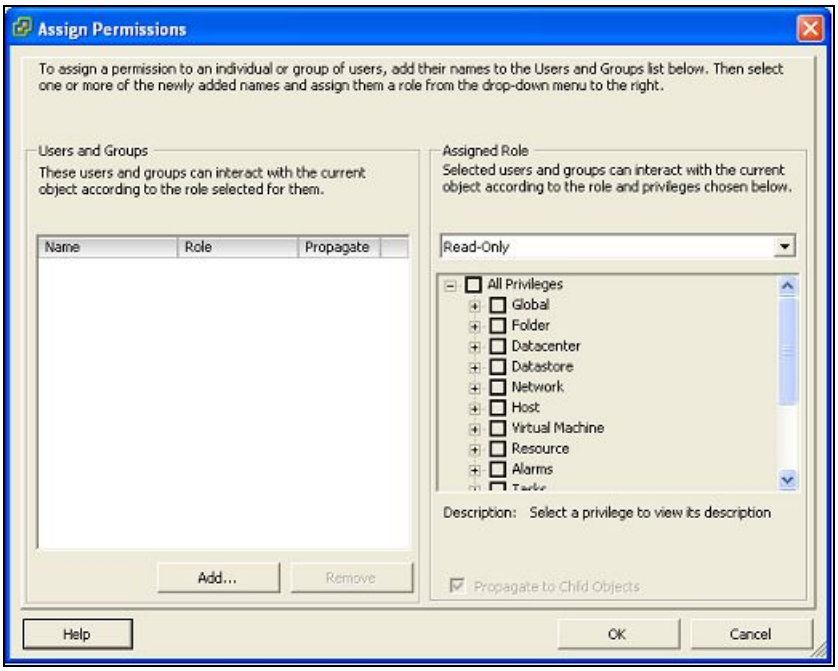

Figure 3.13: Adding a user

11. This will open Figure 3.14. If you want to assign the new role to a domain user, select the **Domain**, and then select the user from the **Users and Groups** list. If you want to assign the new role to a local user, select *(server)* from the **Domain** list, and then select a user from the **Users and Groups** list.

12. Either way, to add the user, click on the **Add** button in Figure 3.14. This will add the chosen user to the **Users** box. Then, click the **OK** button in Figure 3.14.

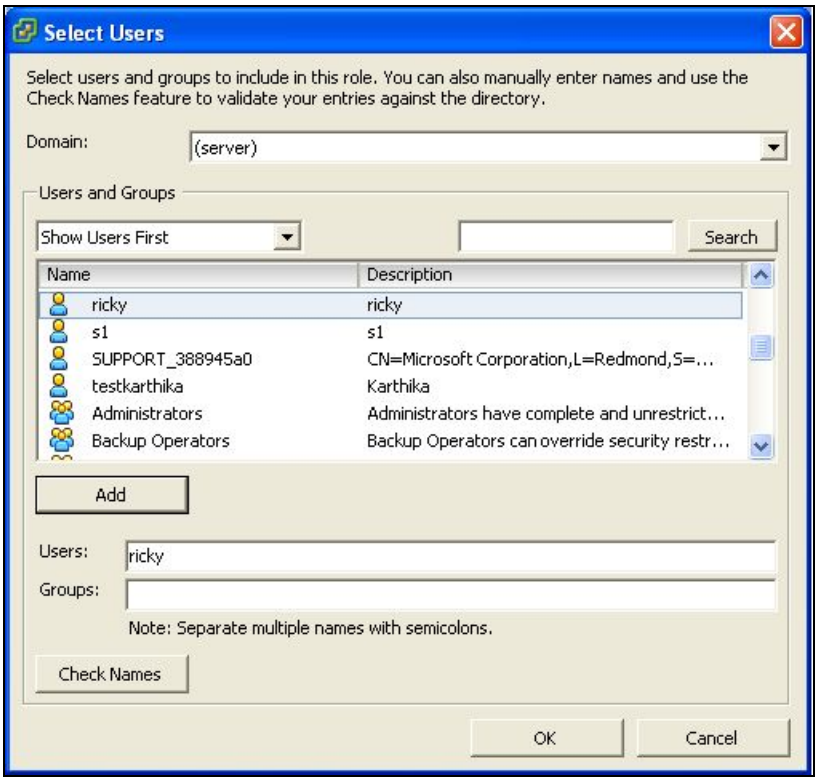

Figure 3.14: Selecting a local/domain user

13. When you return to Figure 3.13, you will find that the local/domain user you selected from Figure 3.14 is displayed in the **Users and Groups** section (see Figure 3.15).

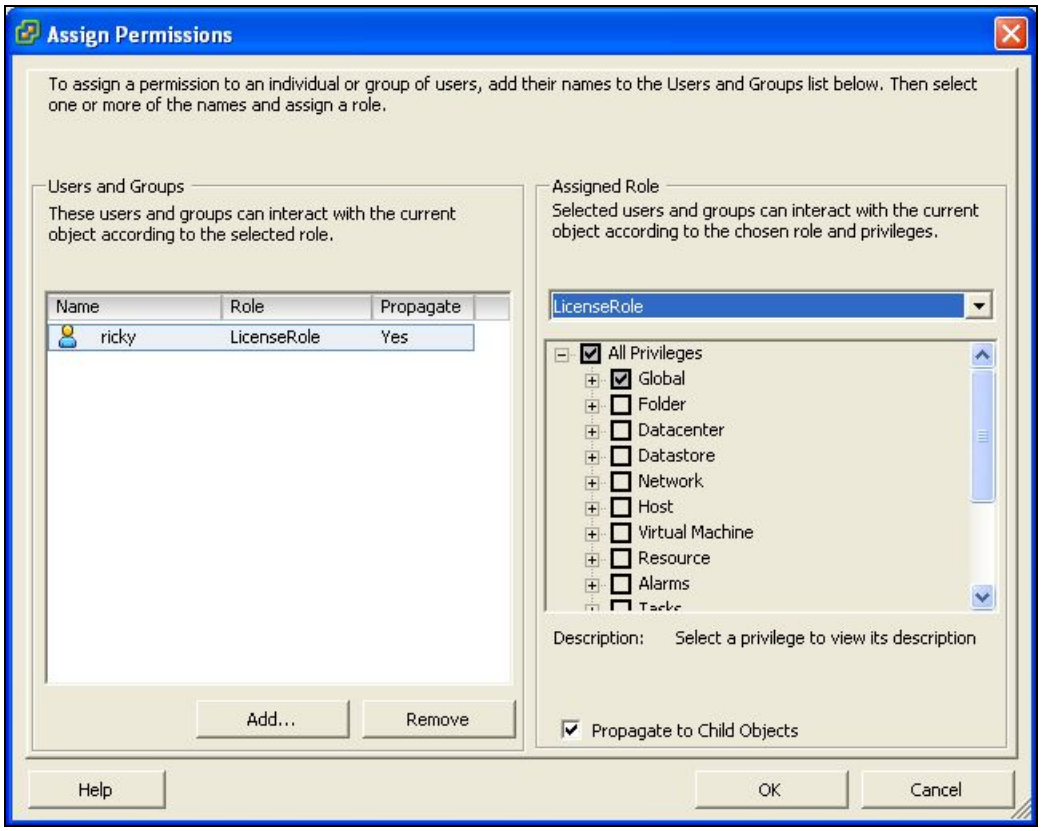

Figure 3.15: Assign the newly created role to the user

14. To assign the newly created role to this user, select the user displayed in the **Users and Groups** section in Figure 3.15, pick the newly created role from the **Assigned Role** section, and then, click the **OK** button in Figure 3.15.
15. Figure 3.16 then appears displaying the user-role mapping.

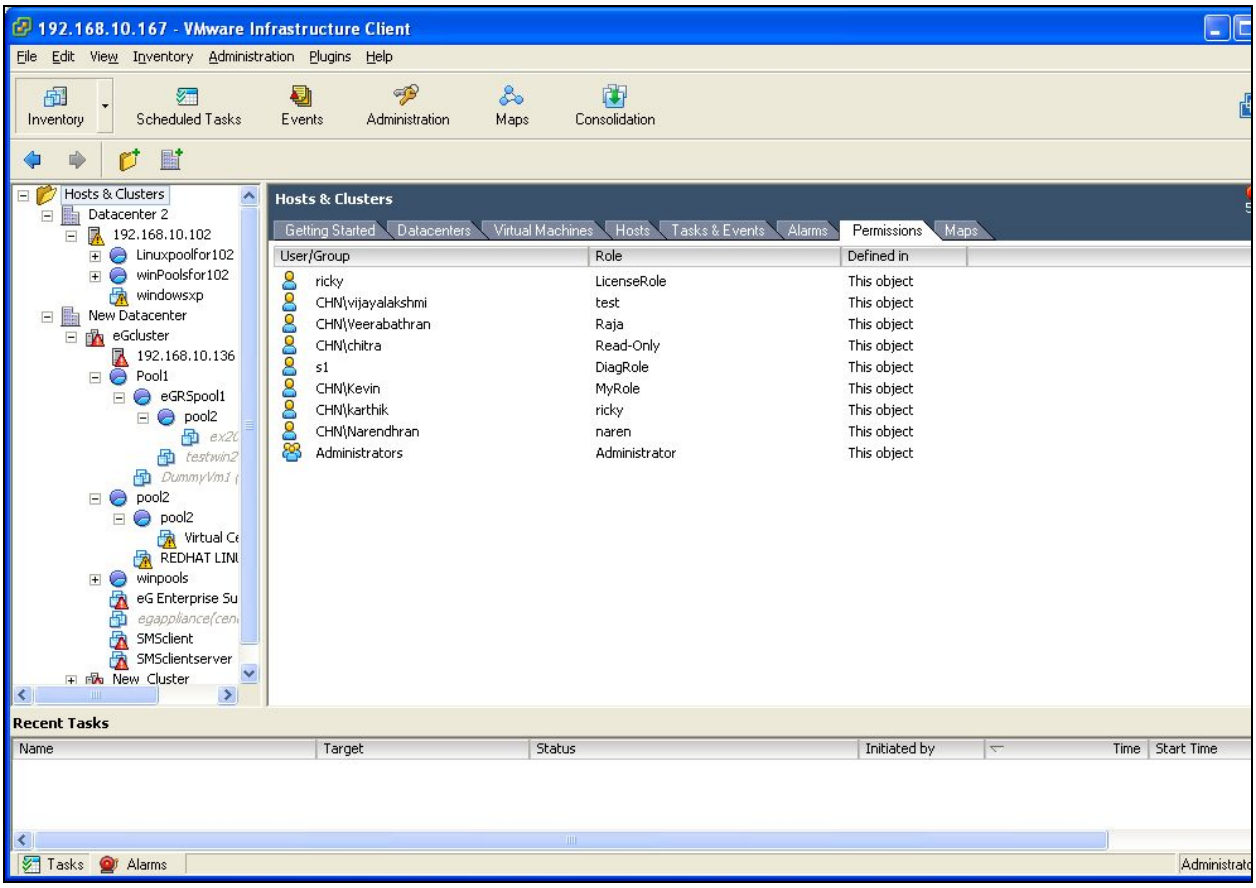

Figure 3.16: The Permissions tab page displaying the chosen user and the role assigned to him/her

3.1.6 ESX Servers Test

This test reports the number of ESX hosts that are currently managed by VC, and the status of the ESX-VC connections.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the vCenter server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.

Parameter	Description
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the SSL flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for

Parameter	Description
	collecting metrics from vCenter.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
ESX servers managed by VCenter	Indicates the number of ESX servers that are currently managed by VC.	Number	Use the detailed diagnosis of this measure, if enabled, to know which ESX servers are being managed by VC.
ESX servers added to VCenter	Indicates the number of ESX servers that were added to VC during this measurement period.	Number	Use the detailed diagnosis of this measure, if enabled, to know which ESX servers were recently added to VC.
ESX servers removed from VCenter	Indicates the number of ESX servers that were removed from VC during this measurement period.	Number	Use the detailed diagnosis of this measure, if enabled, to know which ESX servers were removed from VC.
ESX servers connected to VCenter	Indicates the number of ESX servers that are currently connected to VC.	Number	Use the detailed diagnosis of this measure, if enabled, to know which ESX servers are connected to VC.
ESX servers recently connected to VCenter	Indicates the number of ESX servers that connected to VC during this measurement period.	Number	Use the detailed diagnosis of this measure, if enabled, to know which ESX servers recently connected to VC.

Measurement	Description	Measurement Unit	Interpretation
ESX servers disconnected from VCenter	Indicates the number of ESX servers that are currently disconnected from VC.	Number	Use the detailed diagnosis of this measure, if enabled, to know which ESX servers disconnected from VC.
ESX servers recently disconnected from VCenter	Indicates the number of ESX servers that disconnected from VC during this measurement period.	Number	Use the detailed diagnosis of this measure, if enabled, to know which ESX servers recently disconnected from VC.
ESX servers not responding to VCenter	Indicates the number of ESX servers that are currently not responding to VC.	Number	Ideally, the value of this measure should be low. A high value could indicate a problem in the network connection between the ESX server and VC. Use the detailed diagnosis of this measure, if enabled, to know which ESX servers are not responding to VC.
ESX servers recently not responding from vCenter	Indicates the number of ESX servers that have been marked as 'unresponsive' during this measurement period.	Number	Use the detailed diagnosis of this measure, if enabled, to know which ESX servers were marked as 'not responding' during this measurement period.
Virtual machines managed by VCenter	Indicates the number of VMs that are currently executing on all ESX servers managed by this VC.	Number	Use the detailed diagnosis of this measure to view the names of the VMs, the names of the ESX hosts on which the VMs have been configured, and the IP address of the hosts,.
VM templates managed by VCenter	Indicates the total number of template VMs that currently exist on all ESX servers managed by VC.	Number	A template is a "golden" copy of a virtual machine (VM) organized by folders and managed with permissions. They are useful because they act as a protected version of a model VM which can be used to create new VMs. As a template is the original and perfect image of a particular VM, it cannot be powered on or run.

Measurement	Description	Measurement Unit	Interpretation
			You can use the detailed diagnosis of this measure to view the names and IP addresses of the template VMs.
Orphaned VMs managed by VCenter	Indicates the total number of orphaned vms that have been currently detected on all ESX servers managed by VC.	Number	An orphan virtual machine is one that exists in the vCenter database but is no longer present on the ESX Server host. A virtual machine also shows as orphaned if it exists on a different ESX Server host than the ESX Server host expected by vCenter. A virtual machine can become orphaned, in any of the following situations: • After a VMotion or VMware DRS Migration; • After a VMware HA host failure occurs or after the ESX Server host comes out of maintenance mode; • If you delete a virtual machine outside of vCenter – say, through the VMware Management Interface while VC is down, or through the Virtual Infrastructure (VI) client directly connected to an ESX server host; • If vCenter is restarted while a migration is in progress; this is a temporary situation though; • If you schedule too many virtual machines to be relocated at the

Measurement	Description	Measurement Unit	Interpretation
			same time; To know the names of the orphaned VMs, use the detailed diagnosis of this measure, if enabled.
Disconnected VMs managed by vCenter	Indicates the number of VMs that are currently disconnected from the ESX server.	Number	
Invalid VMs managed by vCenter	Indicates the number of VMs that are currently invalid.	Number	
Inaccessible VMs managed by vCenter	Indicates the number of VMs that are currently inaccessible.	Number	When an ESX Server machine is rebooted or a host agent is restarted, it needs to reload the host agent configuration of each registered virtual machine. If the .vmx file is inaccessible, ESX Server is unable to read the configured name of the virtual machine, and it defaults to "Unknown VM." This is a problem only during restarts. Temporarily losing access to storage does not cause a virtual machine's name to be set to "Unknown VM". The workaround is to rename the virtual machines that have gotten into this state, after they become available again.

The detailed diagnosis of the *ESX servers managed by VC* measure reports the IP and host name of the ESX servers that are currently managed by VC.

Details of ESX servers managed by the vCenter

Time	HostName	IP
Mar 27, 2009 12:02:45		
	esx102	192.168.10.102
	esx3	192.168.10.136
	esx3i	192.168.10.179

Figure 3.17: The detailed diagnosis of the ESX servers managed by VC measure

The detailed diagnosis of the *ESX servers connected to VC* measure reports the IP and host name of the ESX servers that are currently connected to VC.

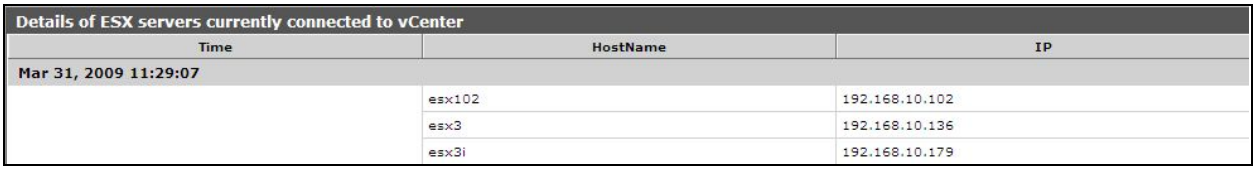

Details of ESX servers currently connected to vCenter

Time	HostName	IP
Mar 31, 2009 11:29:07		
	esx102	192.168.10.102
	esx3	192.168.10.136
	esx3i	192.168.10.179

Figure 3.18: The detailed diagnosis of the ESX servers connected to VC measure

3.1.7 vCenter User Sessions Test

This test reports the number of sessions for every user who is currently logged into VC.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the every currently logged in user to VC.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the

Parameter	Description
	eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the SSL flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current sessions to VCenter	Indicates the number of sessions that this user has currently open on VC.	Number	This measure is a good indicator of the workload generated by a user on VC. The detailed diagnosis of this measure provides the details of every session initiated by this user.

The detailed diagnosis of the *Current sessions to VC* measure reveals the details pertaining to the sessions initiated by a particular user. The details include the name of the user, the login time of the user, and the when last the user session was active on VC.

Details of user sessions to vCenter

Time	UserName	FullName	LoginTime	LastActiveTime
Mar 31, 2009 11:24:07				
	CHN\egtest	eGtest	Tue Mar 31 11:19:34 IST 2009	Tue Mar 31 11:19:51 IST 2009
	CHN\egtest	eGtest	Tue Mar 31 11:19:56 IST 2009	Tue Mar 31 11:19:56 IST 2009
	CHN\egtest	eGtest	Tue Mar 31 11:14:38 IST 2009	Tue Mar 31 11:14:38 IST 2009
	CHN\egtest	eGtest	Tue Mar 31 11:22:40 IST 2009	Tue Mar 31 11:23:45 IST 2009
	CHN\egtest	eGtest	Tue Mar 31 11:23:40 IST 2009	Tue Mar 31 11:23:40 IST 2009
	CHN\egtest	eGtest	Tue Mar 31 11:23:47 IST 2009	Tue Mar 31 11:23:47 IST 2009

Figure 3.19: The detailed diagnosis of the Current sessions to VC measure

3.1.8 vCenter Tasks Test

This test monitors configured VC tasks and reports their current status.

By default, the test monitors a pre-defined set of tasks belonging to pre-configured event categories. You can override this default setting by including/excluding specific tasks from an event category. To achieve this, do the following:

1. Login to the eG administrative interface.
2. Follow the menu sequence: Agents -> Settings -> vCenter Tasks. Alternatively, you can also click on the **Click here** hyperlink that is available above the parameters of this test in the test configuration page.
3. Either way, **vCENTER TASKS** page appears.

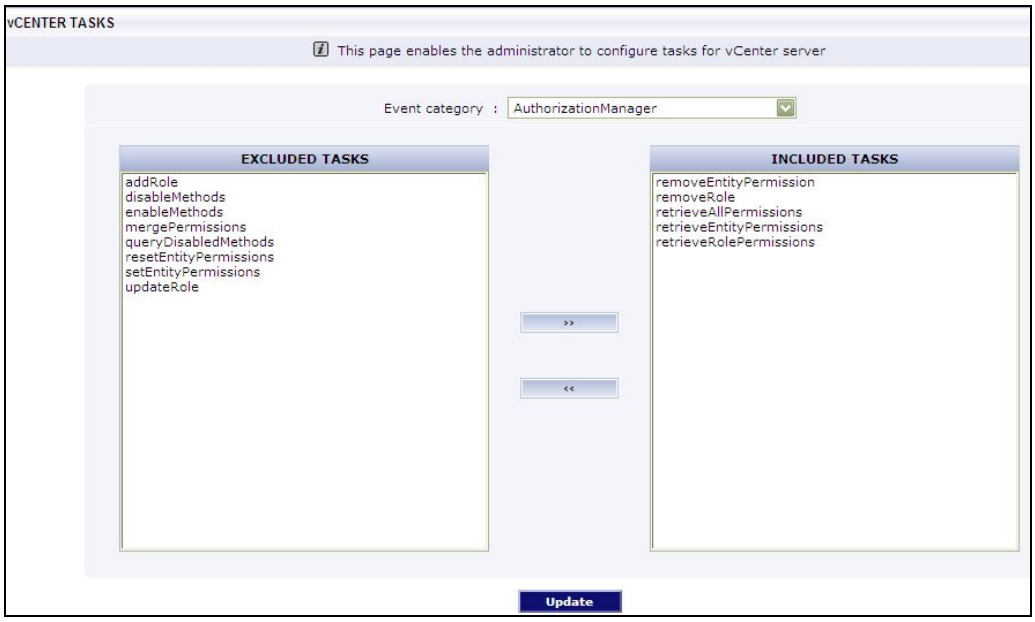

Figure 3.20: Configuring the tasks to be monitored

4. Select an **Event category** to be monitored.

   Note:

   By default, the **Event category** list in the **vCENTER TASKS** page displays a pre-configured set of event categories. You can make more event categories available for selection in this list or remove one/more of the pre-defined categories, if need be. To achieve this, follow the steps given below:

   - Edit the **eg_tests.ini** file in the <EG_INSTALL_DIR>\manager\config directory.
   - In the **[VirtualCenterTasks]** section of this file, you will find the default event categories that will populate the **Event category** list of Figure 26, and the individual events of each category. The format of these default entries is:

     VCTaskEvtTest:<EventCategory>=<Comma- separated list of events belonging to this category>

   - To add a new event category to this default list, you just need to append a line to the **[VirtualCenterTasks]** section in the format mentioned above. For instance, say, you want to include a category named **NfcService** to the **Event category** list of Figure 26. In this case, you will have to insert a line in the **[VirtualCenterEvents]** section for this category, as mentioned below:

     VCTaskEvtTest:NfcService=NfcService.randomAccessOpen,NfcService.randomAccessOpenReadonly,

NfcService.getVmFiles,NfcService.putVmFiles,NfcService.fileManagement,NfcService.system Management

- Next, proceed to the **[VCTaskFilterTypes]** section of the **eg_tests.ini** file. For every event category that is defined in the **[VirtualCenterETasks]** section, a corresponding entry should exist in the **[VCTaskFilterTypes]** section. If not – i.e., if an event category exists only in the [VirtualCenterTasks] section and not in the **[VCTaskFilterTypes]** section – then, such an event category will not be listed as an option in the '**Event Category**' list box.
- Typically, the **[VCTaskFilterTypes]** section is used to indicate which tasks of a particular category need to be monitored by default, and which are not to be monitored. The entries to this section need to be of the following format:

  VCTaskEvtTest:<EventCategory>=<Comma- separated list of tasks to be included while monitoring>:<Comma-separated list of tasks to be excluded from monitoring>

  To include/exclude all tasks of a particular category, use the keyword *all*, and to ensure that no events of a certain category is included/excluded, use the keyword *none*.

- For instance, to ensure that all the tasks of the **Alarms** event category, except the **alarm.Alarm.setCustomValue** task, is to be monitored by default, your specification would be:

  VCTaskEvtTest:Alarms=all:alarm.Alarm.setCustomValue

- Once entries for the new event category are inserted in both the sections, save the **eg_ tests.ini** file.

5. Selecting an **Event category** displays the tasks of the chosen category that are by default included for monitoring in the included tasks list, and those that are by default excluded from monitoring in the excluded tasks list.

   Note:

   - While using the **[VCTaskFilterTypes]** section to configure specific tasks to be included/excluded by default for an event category, make sure that such tasks are a sub-set of the tasks that have been defined for that category in the **[VirtualCenterTasks]** section. For instance, say, you do not want to monitor the **alarm.Alarm.setCustomValue** task of the **Alarms** category. To ensure this, you will have to include the following line in the **[VCTaskFilterTypes]** section:

     VCTaskEvtTest:Alarms=all:alarm.Alarm.setCustomValue

However, before adding this line to the **[VCTaskFilterTypes]** section, verify whether the event, **alarm.Alarm.setCustomValue**,is available in the list of tasks that have been defined for the **Alarms** category in the **[VirtualCenterTasks]** section.

- Make sure that entries such as the following are not inserted in the **[VCTaskFilterTypes]** section:

  VCTaskEvtTest:<EventCategory>=all:all

  VCTaskEvtTest:<EventCategory>=none:none

Also, note that the keywords all and/or none should not be part of a comma-separated list.

Note:

- Every event category that is defined in the **[VirtualCenterTasks]** section of the **eg_tests.ini file** (in the <EG_INSTALL_DIR>\manager\config directory) is associated with a set of tasks. Typically, to indicate which of these tasks need to be monitored by default, and which are not to be monitored, the **[VCTaskFilterTypes]** section of the **eg_tests.ini** file is used.
- For every event category defined in the **[VirtualTaskEvents]** section, a corresponding entry should exist in the **[VCTaskFilterTypes]** section indicating the default set of tasks to be included and/or excluded from monitoring. This entry should be of the following format:

  VCTaskEvtTest:<EventCategory>=<Comma-separated list of tasks to be included while monitoring>:<Comma-separated list of tasks to be excluded from monitoring>

- To include/exclude all tasks of a particular category, use the keyword *all*, and to ensure that no tasks of a certain category is included/excluded, use the keyword *none*.
- For instance, to ensure that all the tasks of the **Alarms** event category, except the **alarm.Alarm.setCustomValue** task, is to be monitored by default, your specification would be:

  VCTaskEvtTest:Alarms=all:alarm.Alarm.setCustomValue

This makes sure that if the **Alarms** option is chosen from the **Event category** list of Figure 26, the **EXCLUDED TASKS** list by default displays the task, **alarm.Alarm.setCustomValue**, and the **INCLUDED TASKS** list displays all other tasks associated with the **Alarms** category.

6. To include any of the tasks that are by default listed in the **EXCLUDED TASKS** list, select the task from the list and click the **<<** button. Similarly, to exclude any of the tasks that are by default

available in the **INCLUDED TASKS** list, select one/more tasks from the list and click the **>>** button.

7. Finally, save the changes by clicking on the **Update** button.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every event category that has been configured for monitoring.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.

Parameter	Description
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
ShowAllTasks	By default, this flag is set to **Yes**. This implies that, by default, this test, in addition to reporting metrics for configured event categories, will also report metrics for an **All** descriptor. Typically, the measures reported by the **All** descriptor will be the aggregate of the measures reported by all the other descriptors of this test – i.e., every measure reported by the **All** descriptor will actually return the sum of the values that all configured tasks have registered for that measure. This enables administrators to easily assess the overall performance of tasks configured for monitoring on a vCenter server.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the SSL flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be

Parameter	Description
	available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Started tasks	Indicates the number of tasks started during this measurement period. For the *All* descriptor, this measure will report the sum of all tasks that started during this measurement period, across all event categories.	Number	
Completed tasks	Indicates the number of tasks that ended during this measurement period. For the *All* descriptor, this measure will report the sum of all completed tasks across all event categories, during this measurement period.	Number	
Successful tasks	Indicates the number of tasks that succeeded during this measurement period. For the *All* descriptor, this measure will report the total number of successful tasks across all event	Number	Use the detailed diagnosis of this measure to view the tasks that succeeded.

Measurement	Description	Measurement Unit	Interpretation
	categories, during this measurement period.		
Failed tasks	Indicates the number of tasks that failed during this measurement period. For the *All* descriptor, this measure will report the sum of all failed tasks across all event categories, during this measurement period.	Number	Ideally, this value should be low. A high value could warrant an investigation into the reason for the consistent failure of tasks. Use the detailed diagnosis of this measure to view the tasks that failed.
Percent of failures	Indicates the percentage of tasks that failed. For the *All* descriptor, this measure will report the sum of the failed percentages registered by all configured tasks under all event categories.	Percent	Ideally, this value should be low. A high value could warrant an investigation into the reason for the consistent failure of tasks.
Avg. completion time	Indicates the average time taken by the tasks of this category to complete. For the *All* descriptor, this measure will report the total time taken by all configured tasks for completion.	Secs	Ideally, this value should be low.
Max completion time	Indicates the maximum time taken by the tasks of this category to complete. For the *All* descriptor, this measure will report the total maximum time taken by all configured tasks under all event categories to complete.	Secs	

Measurement	Description	Measurement Unit	Interpretation
Outstanding tasks	Indicates the number of outstanding tasks of this category during this measurement period. For the *All* descriptor, this measure will report the sum of all outstanding tasks across all event categories during this measurement period.	Number	A large number of outstanding tasks could indicate a bottleneck. Use the detailed diagnosis of this measure to know which are the outstanding tasks.

The detailed diagnosis of the *Successful tasks* measure reveals the name of the successful tasks, the target of the tasks, the user who initiated the tasks, and start time and end time of the tasks.

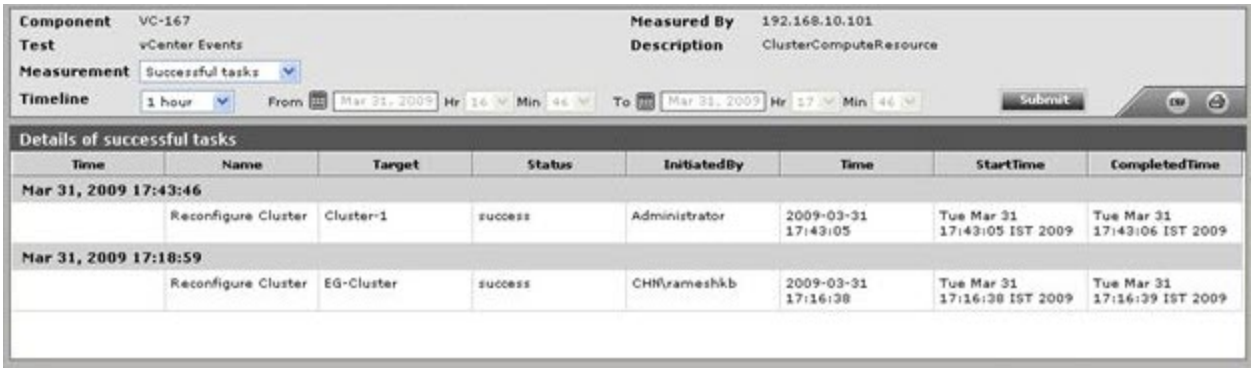

Figure 3.21: The detailed diagnosis of the successful tasks measure

3.1.9 vCenter Events Test

This test enables administrators to promptly capture and report the count and details of critical information, error, warning, and user events that are generated on the vCenter server.

By default, the test monitors a pre-defined set of events belonging to pre-configured event categories. You can override this default setting by including/excluding specific events from an event category. To achieve this, do the following:

1. Login to the eG administrative interface.
2. Follow the menu sequence: Agents -> Settings -> vCenter Events. Alternatively, you can also click on the Click here hyperlink that is available above the parameters of this test in the test configuration page.
3. Either way, **VCENTER EVENTS** page appears.

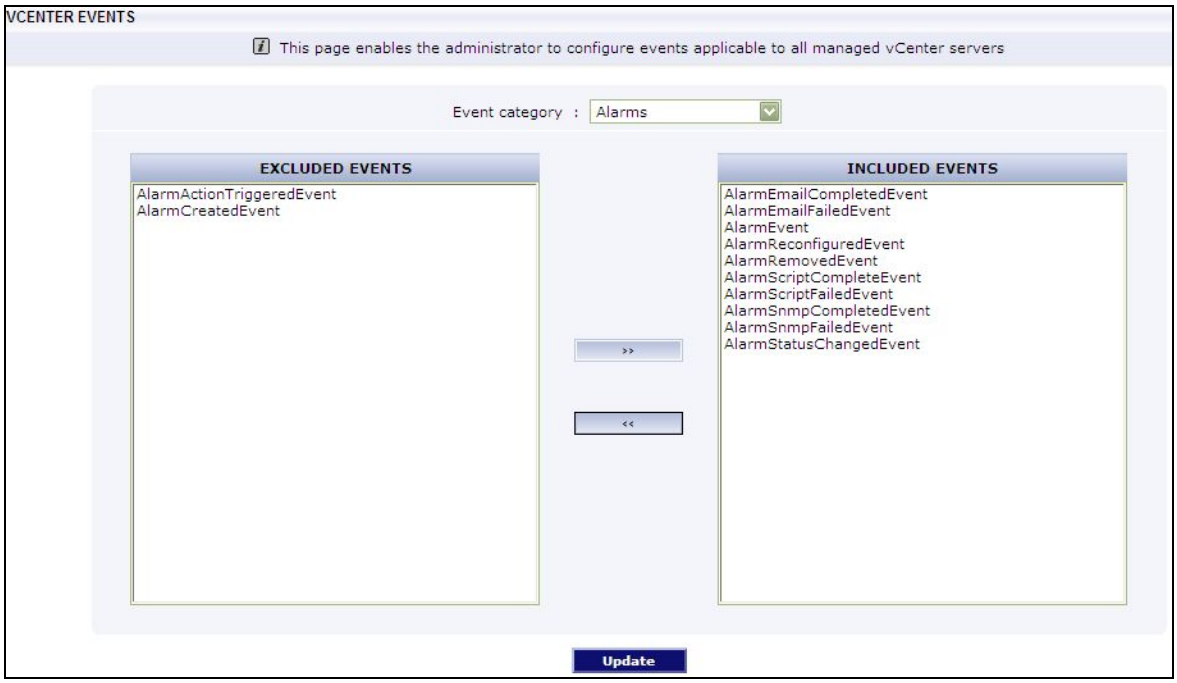

Figure 3.22: Configuring the events to be monitored

4. Select an **Event category** to be monitored.

   Note:

   By default, the **Event category** list in **VCENTER EVENTS** page displays a pre-configured set of event categories. You can make more event categories available for selection in this list or remove one/more of the pre-defined categories, if need be. To achieve this, follow the steps given below:

   - Edit the **eg_tests.ini** file in the <EG_INSTALL_DIRir>\manager\config directory.
   - In the **[VirtualCenterEvents]** section of this file, you will find the default event categories that will populate the **Event category** list and the individual events of each category. The format of these default entries is:

VCEventsTest:<EventCategory>=<Comma- separated list of events belonging to this category>

- To add a new event category to this default list, you just need to append a line to the **[VirtualCenterEvents]** section in the format mentioned above. For instance, say, you want to include a category named **NfcService** to the **Event category** list. In this case, you will have to insert a line in the **[VirtualCenterEvents]** section for this category, as mentioned below:

  VCEventsTest:NfcService=NfcService.randomAccessOpen,NfcService.randomAccessOpenReadonly,NfcService.getVmFiles,NfcService.putVmFiles,NfcService.fileManagement,NfcService.systemManagement

- Next, proceed to the **[VCEventFilterTypes]** section of the **eg_tests.ini** file. For every event category that is defined in the **[VirtualCenterEvents]** section, a corresponding entry should exist in the **[VCEventFilterTypes]** section. **If not – i.e., if an event category exists only in the [VirtualCenterEvents] section and not in the [VCEventFilterTypes] section – then, such an event category will not be listed as an option in the 'Event Category' list box**.

- Typically, the **[VCEventFilterTypes]** section is used to indicate which events of a particular category need to be monitored by default, and which are not to be monitored. The entries to this section need to be of the following format:

  VCEventsTest:<EventCategory>=<Comma-separated list of events to be included while monitoring>:<Comma-separated list of events to be excluded from monitoring>

  To include/exclude all events of a particular category, use the keyword *all*, and to ensure that no events of a certain category is included/excluded, use the keyword *none*.

- For instance, to ensure that all the events of the **Alarms** event category, except the **alarm.Alarm.setCustomValue** event, is to be monitored by default, your specification would be:

  VCEventsTest:Alarms=all:alarm.Alarm.setCustomValue

- Once entries for the new event category are inserted in both the sections, save the **eg_tests.ini** file.

5. Selecting an **Event category** displays the tasks of the chosen category that are by default included for monitoring in the **INCLUDED EVENTS** list, and those that are by default excluded from monitoring in the **EXCLUDED EVENTS** list.

Note:

Every event category that is defined in the **[VirtualCenterEvents]** section of the **eg_tests.ini** file (in the <EG_INSTALL_DIR>\manager\config directory) is associated with a set of events. Typically, to indicate which of these events need to be monitored by default, and which are not to be monitored, the **[VCEventFilterTypes]** section of the **eg_tests.ini** file is used.

For every event category defined in the **[VirtualCenterEvents]** section, a corresponding entry should exist in the **[VCEventFilterTypes]** section indicating the default set of events to be included and/or excluded from monitoring. This entry should be of the following format:

VCEventsTest:<EventCategory>=<Comma-separated list of events to be included while monitoring>:<Comma-separated list of events to be excluded from monitoring>

To include/exclude all events of a particular category, use the keyword all, and to ensure that no events of a certain category is included/excluded, use the keyword *none*.

For instance, to ensure that all the events of the **Alarms** event category, except the **alarm.Alarm.setCustomValue** event, is to be monitored by default, your specification would be:

VCEventsTest:Alarms=all:alarm.Alarm.setCustomValue

This makes sure that if the **Alarms** option is chosen from the **Event category** list, the **EXCLUDED EVENTS** list by default displays the event, **alarm.Alarm.setCustomValue**, and the **INCLUDED EVENTS** list displays all other events associated with the **Alarms** category.

Note:

- While using the **[VCEventFilterTypes]** section to configure specific events to be included/excluded by default for an event category, make sure that such events are a sub-set of the events that have been defined for that category in the **[VirtualCenterEvents]** section. For instance, say, you do not want to monitor the **alarm.Alarm.setCustomValue** event of the **Alarms** category. To ensure this, you will have to include the following line in the **[VCEventFilterTypes]** section:

  VCEventsTest:Alarms=all:alarm.Alarm.setCustomValue

  However, before adding this line to the **[VCEventFilterTypes]** section, verify whether the event, **alarm.Alarm.setCustomValue**,is available in the list of events that have been defined for the **Alarms** category in the **[VirtualCenterEvents]** section.

- Make sure that entries such as the following are not inserted in the **[VCEventFilterTypes]** section:

  VCEventsTest:<EventCategory>=all:all

  VCEventsTest:<EventCategory>=none:none

  Also, note that the keywords *all* and/or *none* should not be part of a comma-separated list.

6. Finally, save the changes by clicking on the **Update** button.

This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick the *VMWare vCenter* as the **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the **<** button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every event category that has been configured for monitoring.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will

Parameter	Description
	remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures*document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Show Information DD	Typically, if the Detailed Diagnosis flag is set to **On** for this test, then periodically, eG Enterprise collects the complete details of all the information, error, warning, and user events generated on vCenter, and stores them in the database. This way, whenever a user clicks on the Diagnosis icon (magnifying glass icon) corresponding to any of the measures reported by this test in the monitoring console, eG Enterprise retrieves the relevant detailed diagnosis information from the database and provides it to the user. In large virtualized environments however, the number of information events generated on the vCenter server will be quite huge. Naturally, the detailed diagnosis of such events will also occupy a considerable amount of database space, which will only grow with time. In order to minimize the strain on the eG database, by default, the detailed diagnosis capability for the information events alone is turned off in the eG Enterprise

Parameter	Description
	system. Accordingly, the Show Information DD flag is set to **No** by default. However, you can this flag is set to **Yes**, so that detailed diagnosis is available for information events as well.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Information	Indicates the number of information events of this category that occurred on vCenter in the last measurement period. For the *All* descriptor, this measure will report the sum of all information events that were captured on vCenter, across all event categories, in the last measurement period.	Number	A change in the value of this measure may indicate infrequent but successful operations performed by one or more events. Use the detailed diagnosis of this measure for more details on the information events.

Measurement	Description	Measurement Unit	Interpretation
Warnings	Indicates the number of warning events of this category that occurred on vCenter in the last measurement period. For the *All* descriptor, this measure will report the sum of all warning events that were captured on vCenter, across all event categories, in the last measurement period.	Number	A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more events. Use the detailed diagnosis of this measure for more details on the warning events.
Errors	Indicates the number of error events of this category that occurred on vCenter in the last measurement period. For the *All* descriptor, this measure will report the sum of all error events that were captured on vCenter, across all event categories, in the last measurement period.	Number	A value (zero) indicates that a vCenter is in a healthy state and is running smoothly without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more events. Use the detailed diagnosis of this measure for more details on the error events.
User events	Indicates the number of user events of this category that occurred on vCenter in the last measurement period. For the *All* descriptor, this measure will report the sum of all user events that were captured on vCenter, across all event categories, in the last measurement period.	Number	A change in the value of this measure may indicate infrequent but successful operations performed by one or more users. Use the detailed diagnosis of this measure for more details on the user events.

The detailed diagnosis of the *Information* measure reveals the description of each information event, the time at which it occurred, the user who initiated the event, and the exact location of the event in terms of the datacenter, cluster, EX host, and VM (if any) on which the event occurred.

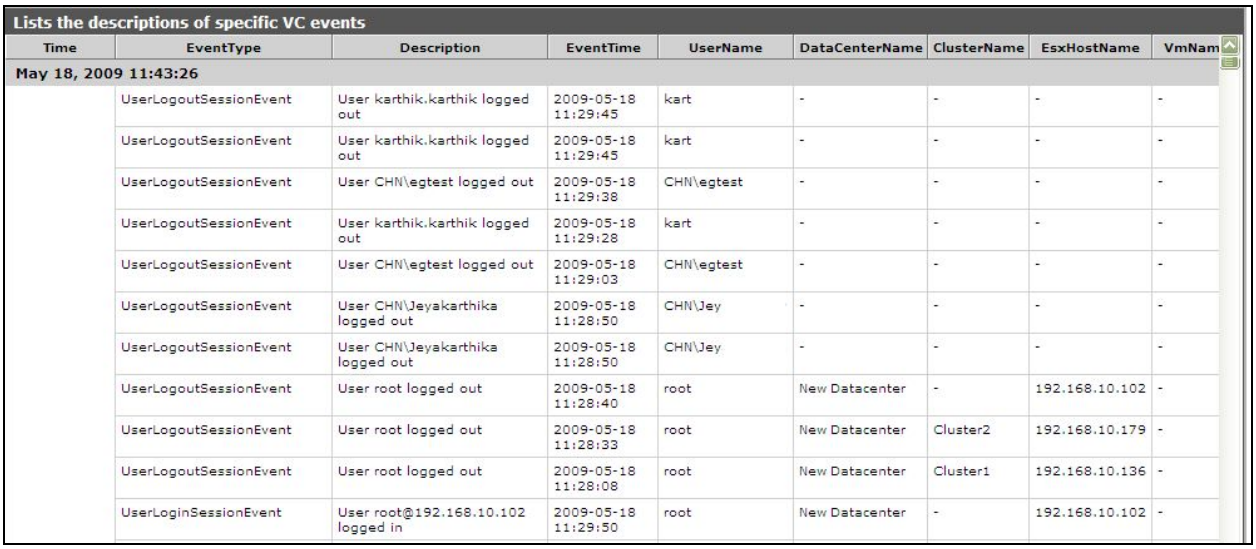

Lists the descriptions of specific VC events

Time	EventType	Description	EventTime	UserName	DataCenterName	ClusterName	EsxHostName	VmNam
May 18, 2009 11:43:26								
	UserLogoutSessionEvent	User karthik.karthik logged out	2009-05-18 11:29:45	kart	-	-	-	-
	UserLogoutSessionEvent	User karthik.karthik logged out	2009-05-18 11:29:45	kart	-	-	-	-
	UserLogoutSessionEvent	User CHN\egtest logged out	2009-05-18 11:29:38	CHN\egtest	-	-	-	-
	UserLogoutSessionEvent	User karthik.karthik logged out	2009-05-18 11:29:28	kart	-	-	-	-
	UserLogoutSessionEvent	User CHN\egtest logged out	2009-05-18 11:29:03	CHN\egtest	-	-	-	-
	UserLogoutSessionEvent	User CHN\Jeyakarthika logged out	2009-05-18 11:28:50	CHN\Jey	-	-	-	-
	UserLogoutSessionEvent	User CHN\Jeyakarthika logged out	2009-05-18 11:28:50	CHN\Jey	-	-	-	-
	UserLogoutSessionEvent	User root logged out	2009-05-18 11:28:40	root	New Datacenter	-	192.168.10.102	-
	UserLogoutSessionEvent	User root logged out	2009-05-18 11:28:33	root	New Datacenter	Cluster2	192.168.10.179	-
	UserLogoutSessionEvent	User root logged out	2009-05-18 11:28:08	root	New Datacenter	Cluster1	192.168.10.136	-
	UserLoginSessionEvent	User root@192.168.10.102 logged in	2009-05-18 11:29:50	root	New Datacenter	-	192.168.10.102	-

Figure 3.23: The detailed diagnosis of the Information measure

The detailed diagnosis of the *Information* measure reveals the description of each error event, the time at which it occurred, the user who initiated the event, and the exact location of the event in terms of the datacenter, cluster, ESX host, and VM (if any) on which the event occurred.

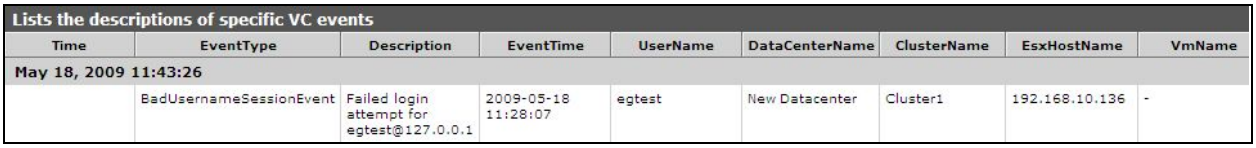

Lists the descriptions of specific VC events

Time	EventType	Description	EventTime	UserName	DataCenterName	ClusterName	EsxHostName	VmName
May 18, 2009 11:43:26								
	BadUsernameSessionEvent	Failed login attempt for egtest@127.0.0.1	2009-05-18 11:28:07	egtest	New Datacenter	Cluster1	192.168.10.136	-

Figure 3.24: The detailed diagnosis of the Error measure

3.1.10 vCenter VMotions Test

VMware vCenter is capable of migrating live virtual machines across entirely separate physical servers with VMware Vmotion, thereby making the maintenance of IT environments non-disruptive. The vCenter Vmotion test tracks the different types of virtual machine migrations occurring on vCenter, and reports their success/failure.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results the vCenter server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is

Parameter	Description
	set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total migrations	Indicates the total number of migrations that occurred on the vCenter during the last measurement period.	Number	Use the detailed diagnosis of this measure to view the details of all migrations that occurred during the last measurement period.
Manual migrations	Indicates the number of powered off or suspended virtual machine moved from one physical server to another during the last measurement period.	Number	Use the detailed diagnosis of this measure to view the details of all manual migrations that occurred during the last measurement period.
Hot migrations	Indicates the number of powered on virtual machine moved from one physical server to another during the last measurement period.	Number	Use the detailed diagnosis of this measure to view the details of all hot migrations that occurred during the last measurement period.

Measurement	Description	Measurement Unit	Interpretation
DRS migrations	Indicates the number of DRS recommended virtual machines moved from one physical server to another during the last measurement period.	Number	Use the detailed diagnosis of this measure to view the details of all DRS-recommended migrations that occurred during the last measurement period.
Migration errors	Indicates the number of migration related errors that occurred in vCenter during the last measurement period.	Number	Ideally, this value should be 0. If the value changes to 1 or more, then you can use the detailed diagnosis of the measure to investigate the error(s).
Migration warnings	Indicates the number of migration related warnings that occurred in vCenter during the last measurement period.	Number	Ideally, this value should be 0. If the value changes to 1 or more, then you can use the detailed diagnosis of the measure to investigate the warning(s).
Migration failures	Indicates the number of migration failures that occurred in vCenter during the last measurement period.	Number	Ideally, this value should be 0. If the value changes to 1 or more, then you can use the detailed diagnosis of the measure to investigate the failure.

3.1.11 vSphere vRAM Usage Test

When VMware released vSphere 5 they also implemented a new licensing model that is based on the amount of virtual RAM (vRAM) that is assigned to virtual machines. vRAM is defined as the memory configured to a virtual machine. vSphere 5 licenses are still sold by the CPU socket, but each license comes with a set vRAM entitlement that varies based on the vSphere 5 edition as shown below:

- vSphere Standard – 32GB per CPU socket
- vSphere Enterprise – 64GB per CPU socket
- vSphere Enterprise Plus – 96GB per CPU socket

If the vRAM license entitlement is exceeded you must purchase another full CPU socket license to increase it. In vSphere 5, there are no limits on the amount of physical memory a host can have; instead, the limits are applied to the amount of vRAM assigned to powered on VMs. VMs that are

powered off do not count against the allotment but VMs that are powered on have the full vRAM that is assigned to a VM counted towards the allotment regardless of the amount of vRAM that a guest OS is actually using. This includes all memory used by a VM, both physical host memory and any memory that a VM may be swapping to disk because of memory over-commitment.

This licensing change in vSphere 5 requires that administrators continually micro-manage their vSphere environment by controlling VM sprawl and by enforcing prudent, calculated usage of the memory resources. This test enables these administrators to better manage the vRAM and the vSphere 5 licenses by periodically reporting the vRAM allocation to VMs and their usage.

Note:

This test is applicable to vCenter 5 only.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results the vCenter server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions,

Parameter	Description
	then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the SSL flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Pooled vRAM capacity	Indicates the amount of vRAM capacity currently available in the license.	GB	The vRAM entitlements of VMware vSphere processor licenses are pooled—that is, aggregated—across all CPU licenses managed by a VMware vCenter instance (or multiple linked VMware vCenter instances) to form a total available vRAM capacity.
VMs powered on	Indicates the number of VMs that are currently powered on.	Number	
Current vRAM usage	Indicates the amount of vRAM currently utilized by the powered-on VMs.	GB	If the vRAM license entitlement is exceeded or exhausted - i.e., if the value of these measures grow close to the value of the Pooled vRAM capacity measure - it is an indication that you must purchase another full CPU socket license. To minimize this financial impact, administrators now need to do a much better job of managing their vSphere environments. No longer can they afford to waste resources and over-allocate memory to virtual machines; doing so now has a financial impact. Administrators also have to be very careful that all virtual machines are right-sized and are not allocated more memory than they actually need to support their workloads. They also need to stay on top of VM lifecycles to ensure that un-needed VMs are deleted when they are no longer needed.
vRAM utilization	Indicates the percentage of total vRAM capacity that is currently utilized by the powered on VMs.	Percent	
Total VMs	Indicates the total number of VMs using this license.	Number	
vRAM usage if all VMs were on	Indicates the amount of vRAM that will be used if	GB	These metrics help administrators

Measurement	Description	Measurement Unit	Interpretation
	all the VMs are powered on.		evaluate the vRAM and license requirement of the virtual environment. By understanding how much vRAM would be required to power on and run all VMs, administrators can determine whether or not the current vRAM (pooled) capacity would be sufficient. If the value of these measures are equal to or higher than the Pooled vRAM capacity measure, it indicates that the total vRAM capacity is not sufficient to support all VMs when powered on. You may hence have to decide what needs to be done to ensure that the virtual environment operates at maximum capacity, but with minimal financial impact: • Whether to procure additional vSphere licenses, so that the aggregate vRAM capacity increases, or, • Whether to power on only a few VMs at a time, so that the vRAM availability is not badly impacted;
vRAM usage if all VMs were on	Indicates the percentage of the pooled vRAM capacity that will be utilized if all the VMs were powered on.	Percent	
vSphere hosts	Indicates the total number of vSphere hosts that are currently using this license.	Number	Use the detailed diagnosis of this measure to know which vSphere/ESX hosts are currently using the license.

3.1.12 VM Movement Test

This test promptly alerts administrators to the addition, removal, and migration of VMs in vCenter. With the help of this test, administrators can determine the number of VMs that were recently added and/or removed from vCenter and the count of vMotions that recently occurred on vCenter.

Moreover, using the detailed diagnosis of the test, administrators can also quickly figure out which VMs were newly added/removed/migrated. When a user complains of being unable to access a VM, administrators can use the detailed diagnosis information to determine the cause of this anomaly – was the VM removed from vCenter at around the time of the complaint? or was the VM being migrated to another vSphere/ESX server when the user complained?

This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the **<** button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

Target of the test : A vCenter server that manages vSphere 5 servers

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the vCenter server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop

Parameter	Description
	Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
VMs created	Indicates number of VMs that were added to the vCenter server during the last measurement period.	Number	To know which VMs were added recently, use the detailed diagnosis of this measure. The following information will be available as part of detailed diagnosis: • The name of the VM that was added; • The host to which the VM was added; • The cluster to which the host belongs; • The resource pool to which the VM belongs; • The CPU, memory, and disk space configuration of the VM; • The datastore assigned to the VM;
Removed VMs	Indicates the number of VMs that were removed from vCenter during the last measurement period.	Number	To know which VMs were deleted recently, use the detailed diagnosis of this measure. The following information will be available as part of detailed diagnosis: • The name of the VM that was removed; • The host from which the VM was removed; • The cluster to which the host belongs;

Measurement	Description	Measurement Unit	Interpretation
			• The resource pool from which the VM was removed; • The CPU, memory, and disk space configuration of the removed VM; • The datastore that was assigned to the VM;
Total migrations	Indicates the total number of migrations that occurred on vCenter during the last measurement period.	Number	Use the detailed diagnosis of this measure to view the details of all migrations that occurred during the last measurement period. This information includes: • The name of the VM that was migrated; • The host from which the VM was migrated; • The host to which the VM was migrated; • The cluster from which the VM was migrated; • The cluster to which the VM was migrated; • The resource pool from which the VM was migrated; • The resource pool to which the VM was migrated; • The CPU, memory, and disk space configuration of the migrated VM;

Measurement	Description	Measurement Unit	Interpretation
			• The datastore assigned to the VM;

3.1.13 vSAN Host Health Status Test

Virtual SAN is a radically simple, hypervisor-converged storage solution for virtual machines. It delivers high performance, scale-out storage that is optimized for vSphere virtual infrastructure. It is an enterprise-class storage solution for any virtualized application, including business-critical workloads. Its seamless integration with vSphere and the entire VMware stack makes it the ideal storage platform for virtual machines. Virtual SAN 6.0 can be configured as hybrid or all-flash storage. In a hybrid storage architecture, Virtual SAN pools server-attached capacity devices, in this case magnetic devices, and caching devices, typically SSDs, and PCI-e devices to create a distributed shared datastore that abstracts the storage hardware and provides a Software-Defined Storage tier for virtual machines.

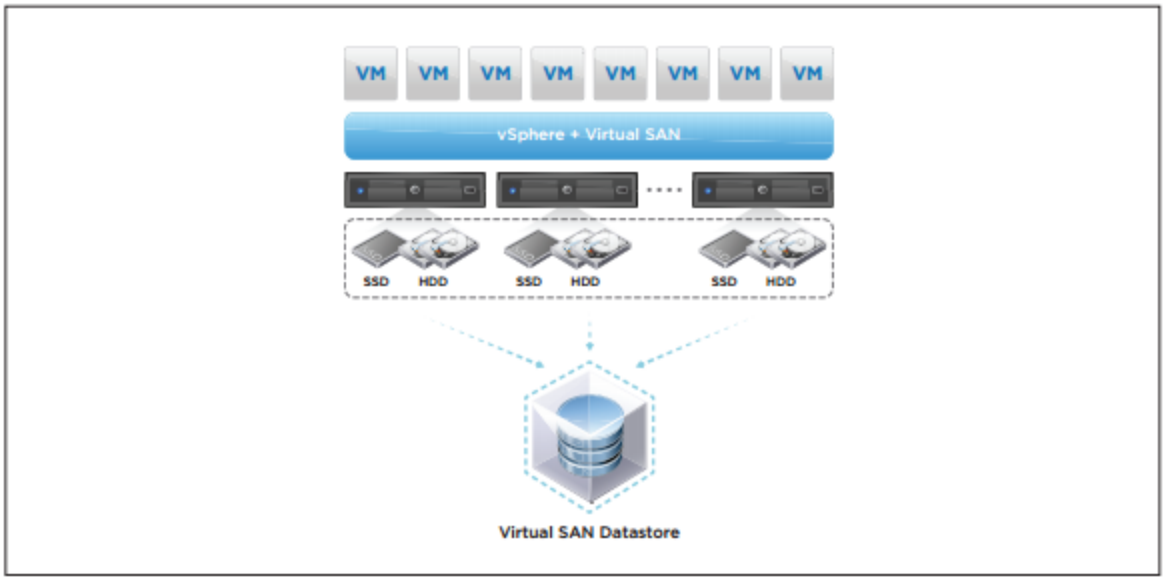

Figure 3.25: VMware virtual SAN

If a host that participates in the vSAN service fails or suffers hardware issues, the vSAN itself may be rendered unavailable/unusable, thus impacting the performance of all the VMs that are using that vSAN. It is hence imperative that administrators monitor the health of every host that contributes its storage resources to the vSAN, quickly identify the host that is experiencing errors/failures, and take pre-emptive measures to fix the faults, so that VM performance is not impacted. This is where the

vSAN Host Health Status test helps. This test alerts administrators to the unhealthy state of any host that is participating in the vSAN service, and ensures their timely intervention in resolving the issues and restoring normalcy.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every vSphere/ESX server host that is participating in the vSAN service.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.

Parameter	Description
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Health status	Indicates whether/not this host is in a healthy state currently.		The values that this measure reports and their corresponding numeric values are listed below: Measure Value: Numeric Value Healthy: 1 Unhealthy: 0 **Note:** By default, this measure reports the **Measure Value**s listed in the table above. In the graph of this measure however, the state of the host is represented using the corresponding numeric equivalents only.

3.2 The Virtual Networking Layer

Using the test mapped to this layer, administrators can figure out the current state of each virtual port and the level of traffic handled by each virtual port.

Figure 3.26: The test mapped to the Virtual Networking layer

3.2.1 Distributed vSwitch Ports Test

A vSphere distributed vSwitch act as a single virtual switch across all associated hosts. This allows virtual machines to maintain a consistent network configuration as they migrate across hosts. A distributed port group is a port group associated with a vSphere distributed switch and specifies port configuration option for each member of port. Distributed port groups define how a connection is made through the vSphere distributed switch in a network. Uplink ports connect a vSphere distributed switch to physical NICs on associated ESX\ESXi hosts. The number of uplinks on a vSphere distributed switch is the maximum number of allowed physical connections to the vSphere distributed switch per host.

This test auto-discovers the virtual ports of the vSphere Distributed Switch and reports the current state of each virtual port. This test also alerts the administrators if a virtual port is blocked and if the DirectPath I/O is active on a virtual port. By continously monitoring the virtual ports, administrators can track the traffic through each virtual port and quickly identify the virtual port that is handling the maximum amount of traffic in terms of multicast packets, broadcast packets etc. Administrators can also figure out the virtual port that is dropping the maximum number of packets upon transmission and reception.

Target of the Test: A VMware vCenter server

Agent deploying the test: An internal agent

Output of the test: One set of results for each *vSphere Distribued Switch:Distributed Virtual Port Group:Virtual Port* of the target VMware vCenter server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.

Parameter	Description
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the SSL flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a

Parameter	Description
	case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measures reported by the test

Measurement	Description	Measurement Unit	Interpretation
Is blocked	Indicates whether/not this virtual port is blocked.	Number	The values reported by this measure and its numeric equivalents are mentioned in the table below: **Measure value** / **Numeric Value**: Yes / 1; No / 0 **Note:** By default, this measure reports the Measure Values listed in the table above to indicate whether/not the virtual port is blocked. The graph of this measure however is represented using the numeric equivalents only - *0* or *1*.
State	Indicates the current state of this virtual port.	Number	The values reported by this measure and its numeric equivalents are mentioned in

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>the table below:
<table><tr><th>Measure value</th><th>Numeric Value</th></tr><tr><td>Link Up</td><td>1</td></tr><tr><td>Link Down</td><td>0</td></tr></table>Note:
By default, this measure reports the Measure Values listed in the table above to indicate the current state of the virtual port. The graph of this measure however is represented using the numeric equivalents only - 0 or 1.
The detailed diagnosis of this measure indicates the name of the entity (server/VM/host) to which the port is connected, MAC Address and vLan ID of the port.</td></tr>
<tr><td>Direct path IO</td><td>Indicates whether the DirectPath I/O to the virtual machine is active on this virtual port.</td><td>Number</td><td>The values reported by this measure and its numeric equivalents are mentioned in the table below:
<table><tr><th>Measure value</th><th>Numeric Value</th></tr><tr><td>Active</td><td>1</td></tr><tr><td>Inactive</td><td>0</td></tr></table>Note:
By default, this measure reports the Measure Values listed in the table above to indicate whether the DirectPath I/O to the virtual machine is active on the virtual port. The graph of this measure however is represented using the numeric equivalents only - 0 or 1.</td></tr>
<tr><td>Broadcast ingress</td><td>Indicates the rate at</td><td>KB/sec</td><td>Compare the value of these measures</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
traffic	which data was received as boadcast packets by this virtual port during the last measurement period.		across the virtual ports to identify the virtual port through which maximum amount of data was sent/received as broadcast packets.
Broadcast egress traffic	Indicates the rate at which data was sent as broadcast packets through this virtual port during the last measurement period.	KB/sec	
Multicast ingress traffic	Indicates the rate at which data was received as multicast packets by this virtual port during the last measurement period.	KB/sec	Compare the value of these measures across the virtual ports to identify the virtual port through which maximum amount of data was sent/received as multicast packets.
Multicast egress traffic	Indicates the rate at which data was sent as multicast packets through this virtual port during the last measurement period.	KB/sec	
Unicast ingress traffic	Indicates the rate at which data was received as unicast packets by this virtual port during the last measurement period.	KB/sec	Compare the value of these measures across the virtual ports to identify the virtual port through which maximum amount of data was sent/received as unicast packets.
Unicast egress traffic	Indicates the rate at which data was sent as unicast packets through this virtual port during the last measurement period.	KB/sec	
Broadcast ingress packets	Indicates the rate at which broadcast	Packets/sec	Compare the value of these measures across the virtual ports to identify the

Measurement	Description	Measurement Unit	Interpretation
	packets were received by this virtual port during the last measurement period.		virtual port through which maximum number of broadcast packets were sent/received.
Broadcast egress packets	Indicates the rate at which broadcast packets were sent through this virtual port during the last measurement period.	Packets/sec	
Multicast ingress packets	Indicates the rate at which multicast packets were received by this virtual port during the last measurement period.	Packets/sec	Compare the value of these measures across the virtual ports to identify the virtual port through which maximum number of multicast packets were sent/received.
Multicast egress packets	Indicates the rate at which multicast packets were sent through this port during the last measurement period.	Packets/sec	
Unicast ingress packets	Indicates the rate at which unicast packets were received by this virtual port during the last measurement period.	Packets/sec	Compare the value of these measures across the virtual ports to identify the virtual port through which maximum number of unicast packets were sent/received.
Unicast egress packets	Indicates the rate at which unicast packets were sent through this port during the last measurement period.	Packets/sec	
Dropped ingress packets	Indicates the rate at which packets were dropped by this port upon reception during	Packets/sec	Ideally, the value of this measure should be zero. Compare the value of this measure across the virtual ports to identify the virtual port that is dropping too many

Measurement	Description	Measurement Unit	Interpretation
	the last measurement period.		packets upon reception.
Dropped egress packets	Indicates the rate at which packets were dropped by this port upon transmission during the last measurement period.	Packets/sec	Ideally, the value of this measure should be zero. Compare the value of this measure across the virtual ports to identify the virtual port that is dropping too many packets upon transmission.
Exception ingress packets	Indicates the rate at which packets that were not expected for e.g., invalid ethertype packets were received by this port during the last measurement period.	Packets/sec	Exceptions is a catch-all bucket for packets that are not expected. For example, you get a packet with vlan tags outside the allowed range, such as packets with an invalid ethernet header. This is slightly different from drops, which can be caused by shaping or a system temporarily running out memory. You may not see many exception packets under normal circumstances. The only place VMware increments this measure is invalid ether type in the packet.
Exception egress packets	Indicates the rate at which packets that were not expected for e.g., invalid ethertype packets were sent through this port during the last measurement period.	Packets/sec	

3.3 The VSAN Layer

This layer appears only if the Virtual SAN is enabled for the clusters in the vCenter. The tests mapped to this layer perform the following functions:

- Report the health of the Virtual SAN enabled cluster
- Monitor the dis groups for congestion and resynchronization caused by failures/maintenance/decommission/object repair, etc.

- Monitor the available resources on the vSAN clusters
- Track the IO operations and throughput of the physical disks on the vSAN enabled clusters

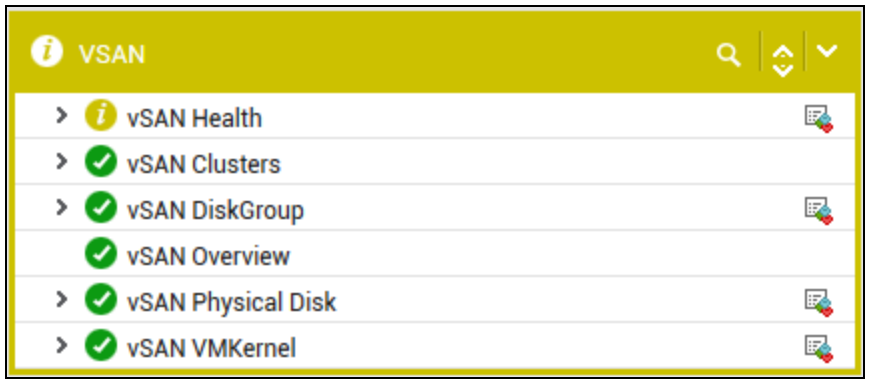

Figure 3.27: The tests mapped to the VSAN layer

3.3.1 vSAN Disk Group Test

vSAN architecture consists of two tiers: a cache tier for the purpose of read caching and write buffering, and a capacity tier for persistent storage. This two tier design offers supreme performance to VMs while ensuring that devices can have data written to them in the most efficient way possible. For this purpose, vSAN uses a logical construct called disk groups to manage the relationship between capacity devices and their cache tier. In a Virtual SAN enabled cluster, multiple disk groups are created on a host so as to improve storage performance significantly while limiting the size of a failure domain; reducing the amount of data impacted by physical device failure.

When a hardware device, host, or network fails, or if a host is placed into maintenance mode or traffic congestion, vSAN initiates resynchronization in the vSAN cluster. However, vSAN might briefly wait for the failed components to come back online before initiating resynchronization tasks. To ensure peak performance even after failures and resynchronization, it becomes necessary for administrators to track the IO operations performed on the disk groups at regular intervals. This is where the **vSAN DiskGroup** test helps administrators.

This test auto-discovers the vSAN disk groups in the vSAN enables clusters in the vCenter and reveals how well the IO operations were performed during resynchronization on each disk group. In the process, this test also reports statistics related to space utilization and congestion on the disk groups. In addition, this test measures the throughput and latency while performing the IO operations in the frontend of each disk group. This revelation sheds light on the IO processing delays, if any, and enables administrators to take necessary actions immediately.

Note:

This test is applicable only for the vSAN enabled clusters in the VMware vCenter server.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each vSAN disk group in the vSAN enabled cluster in the VMware vCenter server.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server

Parameter	Description
	via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Frontend read IOPS	Indicates the number of	IOPS	Virtual machines are considered front-

Measurement	Description	Measurement Unit	Interpretation
	Frontend write operations performed on this vSAN disk group per second.		end – where the application on the virtual machine reads and writes to the vSAN disks. Compare the values of these measures across the disk groups to know what is contributing to the abnormal I/O activity levels - read operations? or write operations?
Frontend write IOPS	Indicates the number of Frontend read operations performed on this vSAN disk group per second.	IOPS	
Frontend read throughput	Indicates the rate at which this disk group processes the Frontend read requests.	MB/Sec	
Frontend write throughput	Indicates the rate at which this disk group processes the Frontend write requests.	MB/Sec	
Frontend read latency	Indicates the average time taken by this disk group to process read Frontend read requests.	Seconds	Ideally, the values of these measures should be very low. By comparing the values of these measures, administrators can figure out where the slowness is maximum - when processing Frontend read requests? or Frontend write requests?
Frontend write latency	Indicates the average time taken by this disk group to process read Frontend write requests.	Seconds	
Capacity	Indicates the total capacity of this vSAN disk group.	GB	
Used capacity	Indicates the amount of space utilized from the total capacity of this vSAN disk group.	GB	Ideally, the value of this measure should be low. If the value of this measure is close to the *Capacity* measure, it indicates that the disk group is running out of space. Compare the value of this measure across the disk groups to identify the disk group that is being over utilized.
Reserved capacity	Indicates the amount of space that has been reserved for future use on this disk group.	GB	The space reserved on each disk group will be provisioned to the hosts after host failures or during maintenance. This way administrators

Measurement	Description	Measurement Unit	Interpretation
			can ensure that sufficient free capacity will be available for components to successfully rebuild after the host failures or during maintenance.
Read Cache hit rate	Indicates the percentage of reads are delivered from the read cache for this disk group.	Percent	A high value is desired for this measure. A grdual/significant decrease in the value of this measure indicates that the read performance is deteriorating while performing read operations from the read cache. In such a case, administrators should increase the size of the vSAN caching tier by adding more disk groups. Alternatively, administrators can tune the working set of the benchmark by doing one of the following: • Decrease the number of active VMs on this cluster • Reduce the number of VM disks accessed by the benchmark • Reduce the size of accessed data in the case of the benchmark
Read cache size	Indicates the size of the read cache managed by this disk group.	GB	VMware vSAN leverage SSD devices of each disk group as the "performance tier" for caching purpose. The purpose of leveraging SSD devices for caching is to serve the highest possible ratio of read operations from the data stored in the read cache and to minimize the read operations to be served by the capacity disks.
Read cache read	Indicates the number of	IOPS	

Measurement	Description	Measurement Unit	Interpretation
IOPS	read operations processed from the read cache.		
Read cache write IOPS	Indicates the number of write operations processed from the read cache.	IOPS	
Read cache Read latency	Indicates the time taken by the read cache for processing the read requests.	Seconds	
Read cache write latency	Indicates the time taken by the read cache for processing the write requests.	Seconds	
Write buffer size	Indicates the size of the write buffer of this disk group.	GB	The vSAN uses "Write buffers" to de-stage written data (not individual write operations) in a way that will create a benign near-sequential (proximal) write workload for the HDDs that form the capacity tier of the vSAN disk group.
Write Buffer free percentage	Indicates the percentage of space available for use in the write buffer of this disk group.	Percent	
Write buffer read IOPS	Indicates the number of read operations processed from the write buffer.	IOPS	
Write buffer write IOPS	Indicates the number of write operations processed from the write buffer.	IOPS	
Write buffer read latency	Indicates the time taken while processing read operations from the write buffer.	Seconds	
Write buffer write latency	Indicates the time taken while processing write	Seconds	

Measurement	Description	Measurement Unit	Interpretation
	operations from the write buffer.		
Bytes De-stage from SSD	Indicates the rate at which the bytes were destaged from the SSD	MB/Sec	
Zero-bytes De-stage		MB/Sec	
Mem congestion	Indicates the number of times the Mem congestion occurred on this disk group.	Number	Congestion is a flow control mechanism used by vSAN. Whenever a bottleneck occurs in a lower layer of vSAN (closer to the physical storage devices), vSAN uses this flow control (aka congestion) mechanism to relieve the bottleneck in the lower layer and instead reduce the rate of incoming I/O at the vSAN ingress, i.e. vSAN Clients (VM Consumption). This reduction of the incoming rate is done by introducing an IO delay at the ingress that is equivalent to the delay the IO would have occurred due to the bottleneck at the lower layer. Thus, it is an effective way to shift latency from the lower layers to the ingress without changing the overall throughput of the system. Mem congestion occurs when the size of used memory heap by vSAN internal components exceed the threshold.
Slab congestion	Indicates the number of times the Slab congestion occurred on this disk group.	Number	Slab congestion is reported when the number of inflight operations exceed the capacity of vSAN internal operation slabs.
SSD congestion	Indicates the number of times the SSD congestion occurred on this disk group.	Number	SSD congestion occurs when the cache tier disk write buffer runs out of space.

Measurement	Description	Measurement Unit	Interpretation
Log congestion	Indicates the number of times the Log congestion occurred on this disk group.	Number	Log congestion occurs when vSAN internal log in cache tier disk runs out of space.
Comp congestion	Indicates the number of times the Comp congestion occurred on this disk group.	Number	Comp congestion occurs when the size of internal table used for vSAN object components is exceeding the threshold.
Cache invalidations	Indicates the number of cache lines that are invalidated due to excessive write operations on this disk group.	Number	Cache invalidations are an indicator for the number of writes on the same address offset as an existing data in the read cache. When a write operation to an IO address follows a read operation, the contents of the read cache must be updated. Such an eviction is referred to as a cache invalidation.
Evictions	Indicates the number of times the read cache contents were evicted due to read cache contention.	Number	Typically, contents in the read cache are evicted when the working set size is larger than the size of the read cache. A low value is desired for this measure. A gradual/sudden increase in the value of this measure indicates the deterioration in the read cache performance.
Outstanding write OPs	Indicates the number of outstanding write operations performed on this disk group.	Number	
Outstanding recovery write OPs	Indicates the number of outstanding recovery write operations performed on this disk group.	Number	
Outstanding write I/O size	Indicates the amount of data written on this disk group during the	GB	

Measurement	Description	Measurement Unit	Interpretation
	outstanding write operations.		
Outstanding recovery write I/O size	Indicates the amount of data written on this disk group during the outstanding recovery write operations.	GB	
Resync read IOPS caused by policy change	Indicates the number of IO read operations used for performing resynchronization on this disk group due to change in policy settings.	IOPS	When there is a change in VM storage policy settings, vSAN might initiate object recreation and subsequent resynchronization of the objects. Compare the values of these measures across the vSAN disk groups to identify the disk group on which maximum number of read and write operations are performed for resynchronization due to change in policy settings.
Resync write IOPS caused by policy change	Indicates the number of IO write operations used for performing resynchronization on this disk group due to change in policy settings.	IOPS	
Resync read IOPS caused by decommission	Indicates the number of IO read operations used for performing resynchronization on this disk group due to decommission.	IOPS	Typically, decommissioning is performed for disk groups from vSAN while upgrading a device or replacing a failed device, or removing a cache device. Compare the values of these measures across the vSAN disk groups to identify the disk group on which maximum number of read and write operations are performed for resynchronization due to decommission.
Resync write IOPS caused by decommission	Indicates the number of IO write operations used for performing resynchronization on this disk group due to decommission.	IOPS	
Resync read IOPS caused by by rebalancing objects	Indicates the number of IO read operations used for performing resynchronization on this disk group while rebalancing the objects.	IOPS	

Measurement	Description	Measurement Unit	Interpretation
Resync write IOPS caused by rebalancing objects	Indicates the number of IO write operations used for performing resynchronization on this disk group while rebalancing the objects.	IOPS	
Resync read IOPS caused by object repair	Indicates the number of IO read operations used for performing resynchronization on this disk group due to the object repair operation.	IOPS	
Resync write IOPS caused by object repair	Indicates the number of IO write operations used for performing resynchronization on this disk group due to the object repair operation.	IOPS	
Resync read throughput caused by policy change	Indicates the rate at which the data is read for performing resynchronization on this disk group due to change in policy settings.	MB/Sec	
Resync write throughput caused by policy change	Indicates the rate at which the data is written for performing resynchronization on this disk group due to due to change in policy settings.	MB/Sec	
Resync read throughput caused by decommission	Indicates the rate at which the data is read for performing resynchronization on this disk group due to the decommission.	MB/Sec	
Resync write	Indicates the rate at which	MB/Sec	

Measurement	Description	Measurement Unit	Interpretation
throughput caused by decommission	the data is written for performing resynchronization on this disk group due to the decommission.		
Resync read throughput caused by rebalancing objects	Indicates the rate at which the data is read for performing resynchronization on this disk group while rebalancing the objects.	MB/Sec	
Resync write throughput caused by rebalancing objects	Indicates the rate at which the data is written for performing resynchronization on this disk group caused by repairing the objects.	MB/Sec	
Resync read throughput caused by object repair	Indicates the rate at which the data is read for performing resynchronization on this disk group caused by repairing the objects.	MB/Sec	
Resync write throughput caused by object repair	Indicates the rate at which the data is written for performing resynchronization on this disk group caused by repairing the objects.	MB/Sec	
Resync read latency caused by policy change	Indicates the average time taken to perform read operations for performing resynchronization due to change in policy settings.	Seconds	vSAN cluster read average latency of resync traffic, including policy change, repair, maintenance mode / evacuation and rebalance from resyncing objects in the perspective of vSAN backend.
Resync write latency caused by	Indicates the average time taken to perform write	Seconds	

Measurement	Description	Measurement Unit	Interpretation
policy change	operations for performing resynchronization due to change in policy settings.		
Resync read latency caused by decommission	Indicates the average time taken to perform read operations during resynchronization due to decommission.	Seconds	
Resync write latency caused by decommission	Indicates the average time taken to perform write operations during resynchronization due to decommission.	Seconds	
Resync read latency caused by rebalancing objects	Indicates the average time taken to perform read operations during resynchronization caused by rebalancing the objects.	Seconds	
Resync write latency caused by rebalancing objects	Indicates the average time taken to perform write operations during performing resynchronization caused by rebalancing the objects.	Seconds	
Resync read latency caused by object repair	Indicates the average time taken to perform read operations during resynchronization due to object repair.	Seconds	
Resync write latency caused by object repair	Indicates the average time taken to perform write operations during resynchronization due to object repair.	Seconds	

3.3.2 vSAN Health Test

In a Virtual SAN enabled cluster, you can use the vSAN health checks to monitor the status of cluster components, diagnose issues, and troubleshoot problems. The health checks cover hardware compatibility, network configuration and operation, advanced vSAN configuration options, storage device health, and virtual machine objects. The vSAN health checks are divided into categories. Each category contains individual health checks.

Health Check Category	Description
Hardware Compatibility	Monitor the cluster components to ensure that they are using supported hardware, software, and drivers.
Performance Service	Monitor the health of vSAN performance service.
Network	Monitor vSAN network health.
Physical disk	Monitor the health of physical devices in the vSAN cluster.
Data	Monitor vSAN data health.
Cluster	Monitor vSAN cluster health.
Capacity utilization	Monitor vSAN cluster capacity.
Online health	Monitor vSAN cluster health and send to VMware's analytics backend system for advanced analysis. You must participate in the Customer Experience Improvement Program to use online health checks.
vSAN Build Recommendation	Monitor vSAN build recommendations for vSphere Update Manager.
vSAN iSCSI target service	Monitor the iSCSI target service, including the network configuration and runtime status.
Encryption	Monitor vSAN encryption health.
Stretched cluster	Monitor the health of a stretched cluster, if applicable.
Hyperconverged cluster configuration compliance	Monitor the status of hosts and settings configured through the Quickstart workflow.

The health checks in the above-table are periodically executed on the vSAN cluster for health testing and performance guarantee. By continuously tracking the health checks on the cluster, administrators can find out current health of the cluster and quickly identify the alerts on unhealthy conditions in time. This way, administrators are enabled to act on the health check alerts that indicate failure conditions or hardware incompatibility with the highest priority. To help administrators in this regard, eG Enterprise offers the **vSAN Health** test.

This test monitors the tests under all the health check categories on the vSAN cluster and reports the count of tests in each health check category at different states. The revelation helps administrators to proactively identify the failures and warnings during health checks and reduces the pain involved in troubleshooting the failure conditions.

Note:

This test is applicable only for the vSAN enabled clusters in the VMware vCenter server.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each *vSAN cluster:health check category* combination.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the

Parameter	Description
	new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
DDForPassedandInfo	By default, both this flag is set to **No**, indicating that by default, the test does not generate detailed diagnostic measures for *Passed* and *Info* measures. If you want the test to generate and store detailed measures for the *Passed* and *Info* measures, set the DDForPassedandInfo flag to **Yes**.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

Parameter	Description
	• The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Passed	Indicates the number of tests under this health check category that returned the Passed state.	Number	The detailed diagnosis of this measure, if enabled using DDForPassedandInfo flag, reveals the name of the tests under each health check category that returned the Passed state, detailed message and health status of the tests.
Skipped	Indicates the number of tests under this health check category that returned the Skipped state.	Number	The detailed diagnosis of this measure lists the name of the tests under each health check category that returned the Skipped state, detailed message and health status of the tests.
Info	Indicates the number of tests under this health check category that returned the Info state .	Number	The detailed diagnosis of this measure, if enabled using DDForPassedandInfo flag, reveals the name of the tests under each health check category that returned the Info state, detailed message and health status of the tests.
Warning	Indicates the number of tests under this health check category that returned the Warning state.	Number	The detailed diagnosis of this measure lists the name of the tests under each health check category that returned the Warning state, detailed message and health status of the tests.
Failed	Indicates the number of tests under this health check category that returned the Failed state.	Number	The detailed diagnosis of this measure lists the name of the tests under each health check category that returned the Failed state, detailed message and health status of the tests.
Unknown	Indicates the number of tests under this health	Number	The detailed diagnosis of this measure lists the name of the tests under each

Measurement	Description	Measurement Unit	Interpretation
	check category that returned the Unknown state.		health check category that returned the Unknown state, detailed message and health status of the tests.

3.3.3 vSAN Overview Test

A VMware vCenter server can manage more than one vSAN cluster. A single vSAN cluster consists of any number of physical server hosts from two to 64. Additional hosts can easily be added to an existing cluster in a matter of minutes for additional compute and storage capacity. vSAN aggregates all local capacity devices into a single datastore shared by all hosts in the vSAN cluster. The datastore can be expanded by adding capacity devices or hosts with capacity devices to the cluster.

To plan the needs of the virtual infrastructure and to determine whether/not additional capacity devices need to be added to the vSAN cluster to handle the current and anticipated load, administrators should keep track on the available resources on the vSAN clusters that are managed by the VMware vCenter server. This can be easily achieved using the **vSAN Overview** test. This test reports the count of datastores, vSAN clusters and hosts available in the VMware vCenter server.

Note:

This test is applicable only for the vSAN enabled clusters in the VMware vCenter server.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for VMware vCenter server that is being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC	To connect to vCenter and extract metrics from it, this test should be configured with

Parameter	Description
Password	the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for

Parameter	Description
	this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
No of datastores	Indicates the total number of datastores in the target VMware vCenter.	Number	The detailed diagnosis of this measure reveals the name of each datastore and the name of datacenter to which the datastore belongs to.
No of clusters	Indicates the total number of vSAN cluster in the target VMware vCenter.	Number	The detailed diagnosis of this measure reveals the name of each cluster and the name of datacenter to which the cluster belongs to.
No of hosts	Indicates the total number of hosts in the target VMware vCenter.	Number	The detailed diagnosis of this measure reveals the name of each host and the name of cluster to which the host belongs to.

The detailed diagnosis of the *No of datastores* measure lists the datastores and the datacenters to which the datastores belong to.

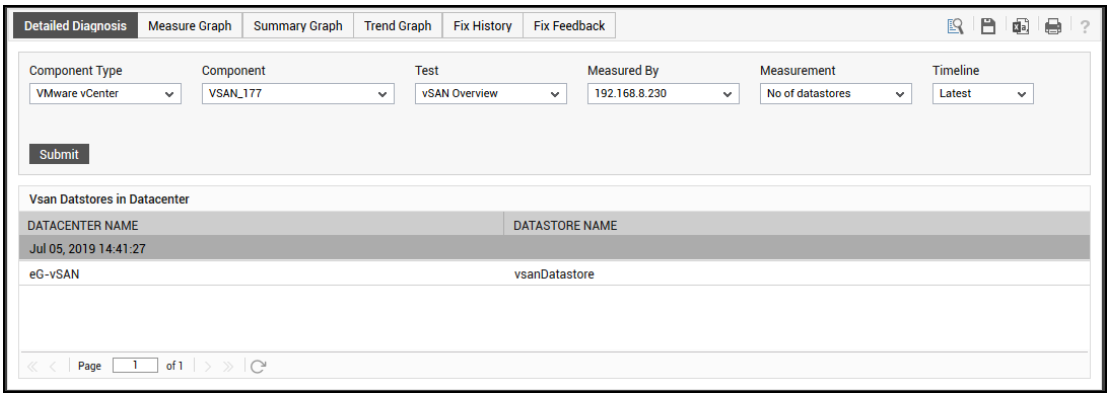

Figure 3.28: The detailed diagnosis of the No of datastores measure

The detailed diagnosis of the *No of clusters* measure lists the clusters and the datacenters to which the clusters belong to.

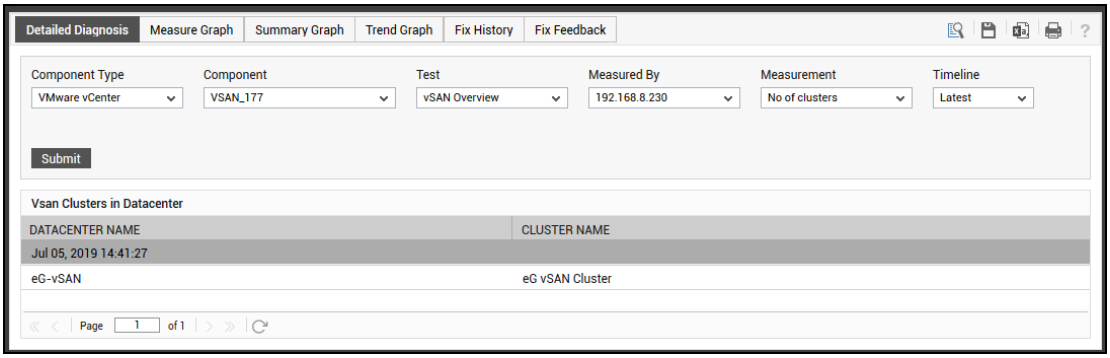

Figure 3.29: The detailed diagnosis of the No of clusters measure

The detailed diagnosis of the *No of hosts* measure lists the hosts and the clusters to which the hosts belong to.

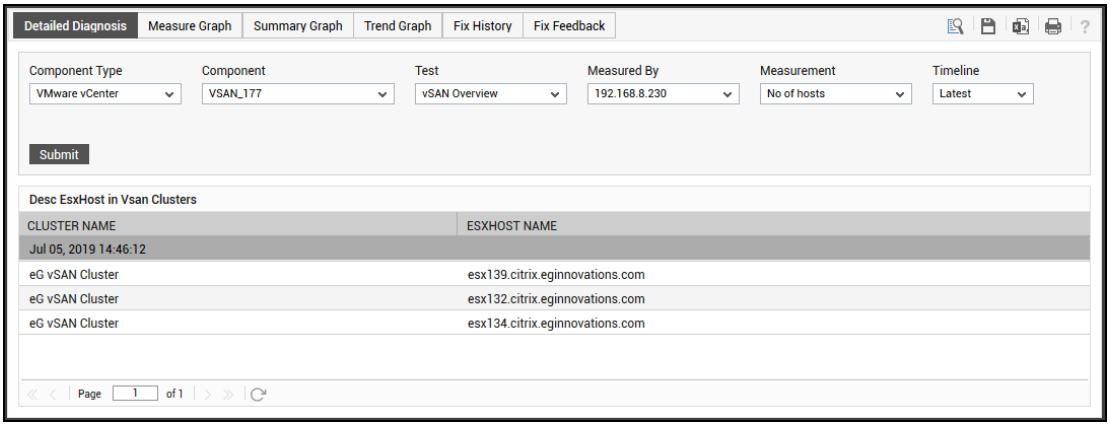

Figure 3.30: The detailed diagnosis of the No of hosts measure

3.3.4 vSAN Physical DiskTest

This test auto-discovers the physical disks in the vSAN cluster and reports the type and current health of each disk. This helps administrators to instantly identify the unhealthy disks and proactively treat the unhealthy disks to prevent prolonged delays in data access for users. This test also reveals the capacity and utilization of each disk, using which any abnormalities can be detected before users start complaining of slowdowns and reduced performance of the cluster. In the process, this test also measures the throughput of read and write operations performed on physical and vSAN layers of each disk. The measured throughput values help administrators to easily find out how well/badly the read and write operations are performed on the physical disks. In addition, the time taken to perform the read and write operations on each disk is also revealed. Using this revelation, administrators can identify the disk which experienced delay while processing the IO operations.

Note:

This test is applicable only for the vSAN enabled clusters in the VMware vCenter server.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for VMware vCenter server that is being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout

Parameter	Description
	period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server,

Parameter	Description
	choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Drive type	Indicates the drive type of this physical disk.		The values this measure reports and their numeric equivalents are provided in the table below: **Measure Value** / **Numeric Value**: FLASH – 0; HDD – 1 **Note:** Typically, this measure reports one of the **Measure Value**s listed in the table above. In the graph of this measure however, the drive type of an physical disk is indicated by its corresponding numeric equivalents only.
Health	Indicates the current health of this physical disk.		The values this measure reports and their numeric equivalents are provided in the table below: **Measure Value** / **Numeric Value**: Healthy – 0; Disk health is unknown – 1

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Permanent disk failure</td><td>2</td></tr>
<tr><td>Permanent disk loss</td><td>3</td></tr>
<tr><td>Disk dicom-missioned</td><td>4</td></tr>
<tr><td>Disk per-formance degraded, and components are evacuating</td><td>5</td></tr>
<tr><td>Disk per-formance degraded, and component evacuation failed</td><td>6</td></tr>
<tr><td>Disk per-formance degraded, and component evacuation get stuck</td><td>7</td></tr>
<tr><td>Disk per-formance degraded, and dying disk is ok to unmount</td><td>8</td></tr>
</table>
Note:
Typically, this measure reports one of the Measure Values listed in the table above. In the graph of this measure however, the health of an physical disk is indicated by its corresponding numeric equivalents only.
</td></tr>
<tr><td>Capacity</td><td>Indicates the total capacity of this disk.</td><td>GB</td><td></td></tr>
<tr><td>Used capacity</td><td>Indicates the amount of space utilized on this disk.</td><td>GB</td><td></td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
Used utilization	Indicates the percentage of space utilized on this disk.	Percentage	
Reserved capacity	Indicates the amount of space that is reserved on this disk for Thick Provisioning.	GB	Some of the objects on vSAN datastore are assigned a storage policy with an Object Space Reservation (OSR) rule set to Thick Provisioning. vSAN reserves the amount of configured capacity for objects with OSR. The capacity is commonly used for an important workload that dynamically consumes storage capacity.
Reserved utilization	Indicates the percentage of space that is reserved on this disk for Thick Provisioning.	Percentage	
Physical layer read IOPS	Indicates the number of read IO operations performed on the Physical layer of this disk.	IOPS	Compare the value of this measure across disks to know which disk handled the maximum number of read requests and which handled the least. If the gap between the two is very high, then it indicates serious irregularities in load-balancing across disks.
Physical layer write IOPS	Indicates the number of write IO operations performed on the Physical layer of this disk.	IOPS	Compare the value of this measure across disks to know which disk handled the maximum number of write requests and which handled the least. If the gap between the two is very high, then it indicates serious irregularities in load-balancing across disks.
Physical layer read throughput	Indicates the rate at which the data was read from the Physical layer of this disk.	MB/sec	A high value is desired for this measure. A very low value is a cause for concern, as it indicates that disk is very poor in handling the read requests.
Physical layer write throughput	Indicates the rate at which the data was written on the Physical	MB/sec	A high value is desired for this measure. A very low value is a cause for concern, as it indicates that disk is very poor in

Measurement	Description	Measurement Unit	Interpretation
	layer of this disk.		handling the write requests.
Physical layer read latency	Indicates the time taken for performing read operations on the Physical layer of this disk.	Seconds	Ideally, this value should be low. If not, it implies that the disk is slow in processing the read requests at the Physical layer.
Physical layer write latency	Indicates the time taken for performing write operations on the Physical layer of this disk.	Seconds	
Guest IO latency	Indicates the time taken for performing IO operations on the guests that share this disk.	Seconds	
Device IO latency	Indicates the time taken for performing IO operations on the devices that share this disk.	Seconds	
vSAN layer read IOPS	Indicates the number of read IO operations performed on the vSAN layer of this disk.	IOPS	
vSAN layer write IOPS	Indicates the number of write IO operations performed on the vSAN layer of this disk.	IOPS	
vSAN layer read latency	Indicates the time taken for performing read operations on the vSAN layer of this disk.	Seconds	
vSAN layer write latency	Indicates the time taken for performing write operations on the vSAN layer of this disk.	Seconds	

3.3.5 vSAN VMKernel Test

A VMkernel network adapter provides network connectivity for hosts and handles the standard system traffic of vSphere vMotion, IP storage, Fault Tolerance, vSAN, and others. Every host that participates in a vSAN cluster should be configured with a VMkernel adapter to handle the vSAN traffic. Smooth traffic flow through the VMkernel adapter is vital for seamless vSAN traffic among the hosts. If high packet loss is monitoring traffic through VMkernel adapter is necessary for administrators to identify the traffic

By continuously monitoring the VMkernel adapters configured on the hosts, administrators can track the traffic through each adapter and quickly identify the virtual port that is handling the maximum amount of traffic in terms of multicast packets, broadcast packets etc. Administrators can figure out the VMkernel adapter that is dropping the maximum number of packets upon transmission and reception. To enable the exchange of data in the vSAN cluster, you must provide a VMkernel network adapter for vSAN traffic on each host.

Note:

This test is applicable only for the vSAN enabled clusters in the VMware vCenter server.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every *host:VMkernel Network Adapter* pair.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been

Parameter	Description
	detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Throughput during reception	Indicates the VMkernel Network Adapter Throughput Inbound.	MB/sec	

Measurement	Description	Measurement Unit	Interpretation
Throughput during transmission	Indicates the VMkernel Network Adapter Throughput Outbound.	MB/sec	
Packets per second during reception	Indicates the VMkernel Network Adapter Inbound Packets Per Second.	Packets/sec	
Packets per second during transmission	Indicates the VMkernel Network Adapter Outbound Packets Per Second.	Packets/sec	
Packets loss during reception	Indicates the percentage of VMkernel Network Adapter Inbound Packets Loss Rate.	Percent	
Packets loss during transmission	Indicates the percentage of VMkernel Network Adapter Outbound Packets Loss Rate.	Percent	

3.4 The Virtual Machine Clusters Layer

This layer reveals how effectively vCenter manages ESX server clusters and cluster resources.

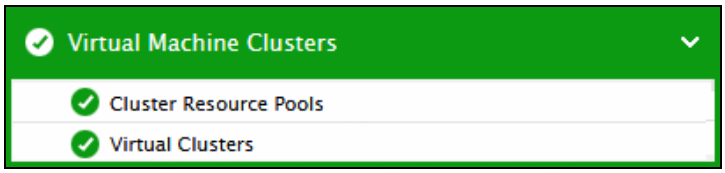

Figure 3.31: The tests mapped to the vCenter Cluster layer

3.4.1 Virtual Clusters Test

This test reports key metrics pertaining to the resource availability and resource usage of the ESX server clusters managed by vCenter.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every *datacenter::clustername* pair managed by vCenter.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag

Parameter	Description
	above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Physical CPU available to cluster	Indicates the aggregated CPU resources of all hosts in the cluster.	Mhz	This measure is the sum total of CPU resources of all hosts in the cluster. The maximum value is equal to the frequency of the processors multiplied by the number of cores.
Effective CPU available to cluster	Indicates the effective CPU resources currently available to run virtual machines.	Mhz	This is the aggregated effective resource level from all running hosts. Hosts that are in maintenance mode or unresponsive are not counted. Resources used by VMware Service console are not included in the aggregate. This value represents the

Measurement	Description	Measurement Unit	Interpretation
			amount of resources available for the root resource pool for running virtual machines.
Physical CPU usage of VMs in cluster	Indicates the current physical CPU usage of the VMs in the cluster in Mhz.	Mhz	
Physical CPU used by VMs	Indicates the percentage of physical CPU resources utilized by the VMs in the cluster.	Percent	Ideally, a cluster should use only a small percentage of the aggregated physical CPU resources of all the hosts within. High CPU usage by one/more hosts in a cluster could drain the physical resources of other hosts in the cluster, thereby affecting the performance of the applications executing on their VMs. In the event of excessive CPU usage by a cluster, it is therefore imperative that you quickly identify which hosts in the cluster, and which VMs on those hosts are responsible for the CPU drain. For this purpose, you can use the detailed diagnosis of this measure, which reveals the VMs in the cluster and the CPU usage of each VM; this way, you can rapidly isolate resource-intensive VMs.
Effective CPU cycles available to the cluster	Indicates the unused CPU currently available with the cluster.	Mhz	This is the difference between the effective CPU resources available with the cluster for running VMs and the effective CPU resources used by the cluster. Ideally, the value of this measure should be high. A very low value or a value that consistently sinks could be a cause for concern, as it could indicate a CPU resource contention on the

Measurement	Description	Measurement Unit	Interpretation
			cluster.
Total physical memory available to cluster	Indicates the aggregated memory resources of all hosts in the cluster.	MB	
Effective memory available to cluster	Indicates the effective memory resources currently available to run virtual machines.	MB	This is the aggregated effective resource level from all running hosts. Hosts that are in maintenance mode or unresponsive are not counted. Resources used by VMware Service console are not included in the aggregate. This value represents the amount of resources available for the root resource pool for running virtual machines.
Consumed memory of host in cluster	Indicates the amount of memory currently used by the hosts in the cluster.	MB	
Physical memory used	Indicates the percentage of physical memory resources utilized by the cluster.	Percent	Ideally, a cluster should use only a small percentage of the aggregated physical memory resources of all the hosts within. High memory usage by one/more hosts in a cluster could drain the physical resources of other hosts in the cluster, thereby affecting the performance of the applications executing on their VMs.
Total CPU cores for cluster	Indicates the number of physical CPU cores currently within the cluster; these are the processes contained by the CPU package.	Number	
Total physical hosts in cluster	Indicates the total number of physical hosts in the cluster, currently.	Number	The detailed diagnosis of this measure, if enabled, reveals the name and IP address of the physical hosts in the

Measurement	Description	Measurement Unit	Interpretation
			cluster.
Number of effective hosts in cluster	Indicates the total number of effective hosts in the cluster, currently.	Number	
VMs in cluster	Indicates the number of VMs in the cluster.	Number	The detailed diagnosis of this measure, if enabled, reveals the name of the VM, the ESX server on which the VM is executing, and and IP address of the ESX server.
VMs powered on	Indicates the total number of VMs that are currently powered-on in the cluster.	Number	The detailed diagnosis of this measure, if enabled, reveals the details of the powered-on VMs.
VMs powered off	Indicates the total number of VMs that are currently powered-off in the cluster.	Number	
VMs suspended	Indicates the number of VMs that are currently in a suspended state in the cluster.	Number	
Effective CPU Reserved	Indicates the total amount of CPU resources that have been used to satisfy the reservation requirements of child resource pool and VMs.	Mhz	
Effective CPU Not Reserved	Indicates the amount of effective CPU resources with the cluster that are currently unused.	Mhz	Ideally, the value of this measure should be high.
CPU Reserved for VMs in Cluster	Indicates the total amount of CPU reserved for VMs in the cluster.	Mhz	Use the detailed diagnosis of this measure to figure out how much CPU has been reserved for each of the VMs in the cluster.
CPU Reserved for Direct Pools in	Indicates the total amount of CPU resources	Mhz	Use the detailed diagnosis of this measure to figure out how much CPU

Measurement	Description	Measurement Unit	Interpretation
Cluster	reserved by the direct pools in the cluster.		has been reserved by each of the direct pools in the cluster.
CPU used by direct Pools in Cluster	Indicates the amount of CPU currently used by direct pools in the cluster.	Mhz	In the event of abnormal CPU usage at the pool-level, use the detailed diagnosis of this measure to identify which direct pool is consuming CPU excessively.
Effective Memory Reserved	Indicates the total amount of memory resources that have been used to satisfy the reservation requirements of child resource pool and VMs.	MB	
Effective Memory Not Reserved	Indicates the amount of effective memory resources with the cluster that are currently unused.	MB	Ideally, the value of this measure should be high.
Memory Reserved for VMs in Cluster	Indicates the total amount of memory reserved for VMs in the cluster.	MB	Use the detailed diagnosis of this measure to figure out how much memory has been reserved for each of the VMs in the cluster.
Memory Reserved for direct Pools in Cluster	Indicates the total amount of memory resources reserved by the direct pools in the cluster.	MB	Use the detailed diagnosis of this measure to figure out how much memory has been reserved by each of the direct pools in the cluster.
Memory consumed for direct Pools in Cluster	Indicates the amount of memory currently used by direct pools in the cluster.	MB	A high value is indicative of excessive memory usage by the direct pools in the cluster. In such circumstances, take the help of the detailed diagnosis capability of this measure to quickly zero-in on that direct pool that consumes maximum memory.
Memory consumed by VMs in Cluster	The total amount of memory currently consumed by the VMs in the cluster.	MB	A high value is indicative of excessive memory usage by the VMs in the cluster. In such circumstances, take the help of the detailed diagnosis

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>capability of this measure to quickly zero-in on that VM that consumes maximum memory.</td></tr>
<tr><td>Total migrations by VMotion</td><td>Indicates the total number of migrations with VMotion that have been done internal to this cluster.</td><td>Number</td><td></td></tr>
<tr><td>HA enabled</td><td>Indicates whether/not this cluster is HA-enabled.</td><td></td><td>If vSphere HA (High Availability) is enabled on a cluster, then vSphere will provide high availability for virtual machines on the cluster. Hosts in the cluster are monitored and in the event of a failure, the virtual machines on a failed host are restarted on alternate hosts.

If vSphere HA is enabled on a cluster, then the value of this measure will be Yes. If not, then this measure will report the value No.

The numeric values that correspond to these measure values are discussed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

By default, this measure reports the Measure Values listed in the table above. In the graph of this measure however, the vSphere HA state is represented using the corresponding numeric equivalents only.</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
DPM enabled	Indicates whether/not the DPM service is enabled for this cluster.		VMware Distributed Power Management (VMware DPM) provides additional power savings by dynamically right-sizing cluster capacity according to workload demands. VMware DPM recommends the evacuation and powering off of ESX hosts when both CPU and memory resources are lightly utilized. VMware DPM recommends powering ESX hosts back on when either CPU or memory resource utilization increases appropriately or additional host resources are needed to meet VMware HA or user-specified constraints. VMware DPM executes VMware DRS in a what-if mode to ensure its host power recommendations are consistent with the cluster constraints and objectives being managed by VMware DRS. Since VMware DPM works in conjunction with VMware DRS, this measure will report a value only if the 'DRS enabled' measure reports the value '*Yes*'. If DPM is enabled for a cluster, then the value of this measure will be *Yes*. If not, then the measure will report the value *No*. The numeric values that correspond to these measure values are discussed in the table below:

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Yes</td><td>1</td></tr>
<tr><td>No</td><td>0</td></tr>
</table>
Note:
By default, this measure reports the Measure Values listed in the table above. In the graph of this measure however, the vSphere DPM state is represented using the corresponding numeric equivalents only.
</td></tr>
<tr><td>DRS enabled</td><td>Indicates whether/not the DRS service is enabled for this cluster.</td><td></td><td>DRS (Dynamic Resource Scheduler) is a powerful feature that enables your virtual environment to automatically balance itself across your ESX host servers in an effort to eliminate resource contention. It utilizes the VMotion feature to provide automated resource optimization through automatic migration of VMs across hosts in a cluster. DRS also provides automatic initial VM placement on any of the hosts in the cluster, and makes automatic resource relocation and optimization decisions as hosts or VMs are added to or removed from the cluster. You can also configure DRS for manual control so that it only provides recommendations that you can review and carry out.
DRS works by utilizing resource pools and clusters that combine the resources of multiple hosts into a single entity. When a VM experiences increased load, DRS first evaluates its priority against the established</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>resource allocation rules and then, if justified, redistributes VMs among the physical servers to try to eliminate contention for resources. VMotion will then handle the live migration of the VM to a different ESX host with complete transparency to end users. The dynamic resource allocation ensures that capacity is preferentially dedicated to the highest-priority applications, while at the same time maximizing overall resource utilization.

If DRS is enabled for a cluster, then the value of this measure will be Yes. If not, then the measure will report the value No.

The numeric values that correspond to these measure values are discussed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

By default, this measure reports the Measure Values listed in the table above. In the graph of this measure however, the vSphere DRS state is represented using the corresponding numeric equivalents only.</td></tr>
<tr><td>vSAN service enabled</td><td>Indicates whether/not the vSAN service is enabled for this cluster.</td><td></td><td>Virtual SAN virtualizes local physical storage resources of vSphere hosts and turns them into pools of storage that can be carved up and assigned to virtual machines and applications</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>according to their quality of service requirements.

You can activate Virtual SAN when you create host clusters or enable Virtual SAN on existing clusters. When enabled, Virtual SAN aggregates all local storage disks available on the hosts into a single datastore shared by all hosts. You can later expand the datastore by adding storage devices or hosts to the cluster.

If vSAN is not enabled for a cluster, then the value of this measure will be No. If vSAN is enabled, then the value of this measure will be Yes. The numeric values that correspond to these measure values are discussed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

By default, this measure reports the Measure Values listed in the table above. In the graph of this measure however, the state of the vSAN is represented using the corresponding numeric equivalents only.</td></tr>
</table>

The detailed diagnosis of *Total physical hosts in cluster* measure, if enabled, reveals the name and IP address of the physical hosts in the cluster.

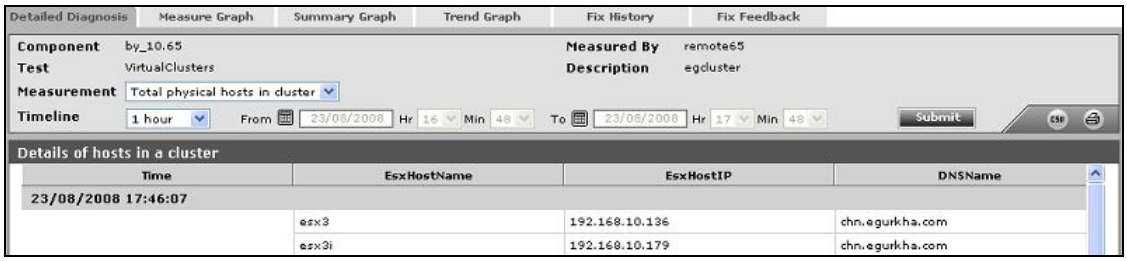

Figure 3.32: Figure 27: The detailed diagnosis of the Total physical hosts in cluster measure

The detailed diagnosis of the *VMs in cluster* measure, if enabled, reveals the name of the VM, the ESX server on which the VM is executing, and and IP address of the ESX server.

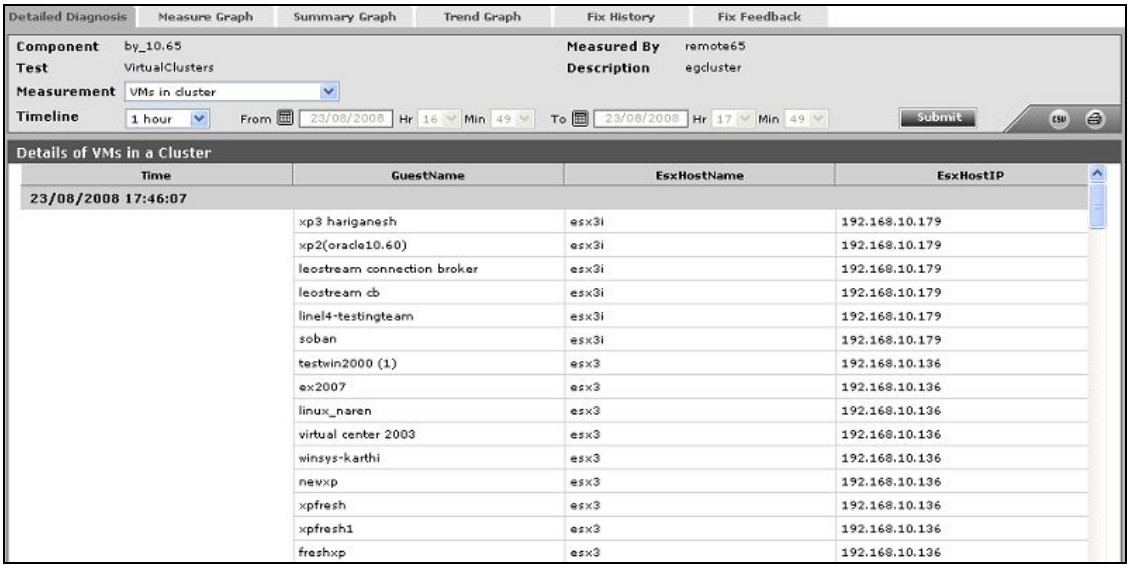

Figure 3.33: The detailed diagnosis of the VMs in cluster measure

The detailed diagnosis the *VMs powered on* measure, if enabled, reveals the details of the powered-on VMs.

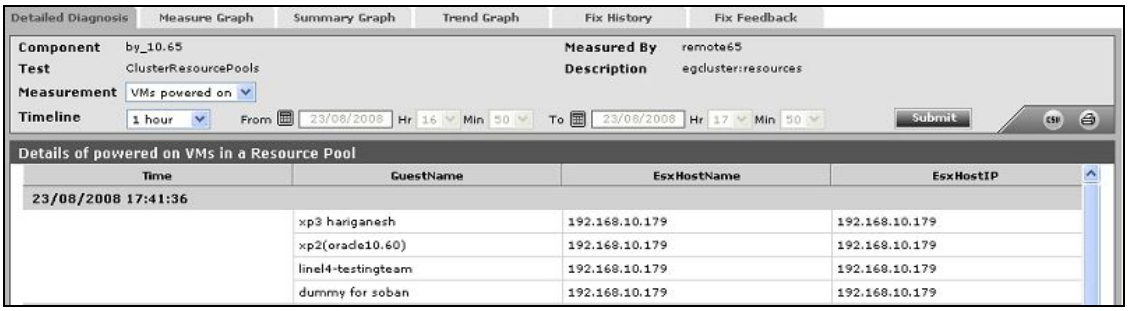

Figure 3.34: The detailed diagnosis of the VMs powered on measure

3.4.2 Cluster Resource Pools Test

This test auto-discovers the resource pools/vApps under ESX clusters, and reports critical statistics pertaining to each pool.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every *datacenter::cluster::resourcepool and/or datacenter::cluster::vApp* combination managed by vCenter; ; to distinguish between vApps and resource pools, the eG monitoring console tags every discovered vApp as [vApp].

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the

Parameter	Description
	Monitoring VMware Infrastructures document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *4:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Pool status</td><td>Indicates the current status of this resource pool.</td><td></td><td>If the status cannot be determined, then Unknown will be displayed here. If the pool is functioning normally, then this measure will return the value Normal. If the root resource pool does not have the capacity to meet the reservation of its children, then the status will be displayed as Overcommitted. One or more pools in the tree may have children whose reservations are greater than the pool is configured to support, then the status will be Inconsistent.

The numeric values that correspond to these measures are as follows:

<table><tr><th>State</th><th>Value</th></tr><tr><td>Unknown</td><td>0</td></tr><tr><td>Normal</td><td>1</td></tr><tr><td>Overcommitted</td><td>2</td></tr><tr><td>Inconsistent</td><td>3</td></tr></table>
Note:

By default, this measure reports the Measure Values listed in the table above to indicate the status of each resource pool. The graph of this measure however, represents the status of a resource pool using the numeric equivalents - 1 to 3.</td></tr>
<tr><td>Child resource pools</td><td>Indicates the total number of child resource pools under this resource pool.</td><td>Number</td><td></td></tr>
<tr><td>VMs in pool</td><td>Indicates the total number of VMs in this resource pool.</td><td>Number</td><td>This number includes both the directly and indirectly assigned VMs.</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
VMs powered on	Indicates the total number of powered-on VMs in this resource pool.	Number	This number includes both the directly and indirectly assigned VMs.
Direct VMs in pool	Indicates the number of virtual machines that were directly assigned to this resource pool.	Number	
Direct VMs powered on	Indicates the number of directly assigned virtual machines under this resource pool that is currently powered on.	Number	If the value of this measure is equal to the value of the Direct VMs in pool measure, it indicates that all VMs in the pool are in a powered-on state. If the values do not match, it indicates that one/more VMs are currently in a powered-off state.
CPU reserved for VMs in pool	The total amount of CPU reserved by VMs in this pool.	Mhz	Use the detailed diagnosis of this measure to know how much CPU has been reserved by each VM.
Host CPU used by VMs	Indicates the percentage of physical CPU resources used by the VMs in this pool.	Percent	A high value of this measure could indicate excessive CPU usage by the child resource pools/virtual machines configured under this resource pool. By comparing the value of this measure across resource pools, you can quickly identify that resource pool which is draining the CPU resources of the ESX host. The detailed diagnosis of this measure, if enabled, reveals the names and CPU usage of each VM in this pool, thus revealing which VM is CPU-intensive.
CPU capacity used	Indicates the CPU usage of the resource pool in Mhz	Mhz	The percent CPU usage measure serves as an effective indicator of how resource-intensive a particular resource pool is on a specific ESX server. However, for performing capacity planning or what-if analysis, the CPU

Measurement	Description	Measurement Unit	Interpretation
			usage of a pool measured in absolute terms would be more useful. For instance, the physical CPU usage of a pool could be 30% on a particular ESX server – this means that the resource pool is consuming 30% of the total physical CPU capacity of that ESX server. If you are planning to migrate the resource pool to another ESX server, then it would be unwise to assume that the pool will only consume 30% of CPU on the other ESX server as well, as the percentage will vary depending upon the physical CPU resources that are available to the other ESX server. The absolute measure however will remain unchanged across ESX server. Therefore, to decide which ESX server a VM/resource pool is to be moved to, and to analyze the impact of this movement on the CPU resources of the new ESX host, you would require an absolute measure of CPU usage.
CPU capacity in pool	Indicates the amount of CPU resources that is guaranteed available to the virtual machine or resource pool.	Mhz	Typically, the Reservation setting for a resource pool specifies the minimum acceptable amount of CPU or memory - not the amount you would like to have available. Follow the guidelines mentioned below to define the Reservation setting: • Do not set the Reservation too high. A very high reservation can limit the number of virtual machines you can power on in a resource pool. • Always leave some headroom while reserving resource. Do not commit all

Measurement	Description	Measurement Unit	Interpretation
			resources. As you move closer to fully reserving all capacity in the system, it becomes increasingly difficult to make changes to reservations and to the resource pool hierarchy without violating admission control. • If you expect frequent changes to the total available resources, use Shares, not Reservation, to allocate resources fairly across resource pools. Reserved resources are not wasted if they are not used. If the utilization is less than the reservation, the resources can be utilized by other running virtual machines.
CPU capacity reserved	Indicates the total amount of CPU resources that have been used to satisfy the reservation requirements of all descendants of this resource pool (includes both resource pools and virtual machines).	Mhz	CPU resources are typically reserved for powered-on virtual machines in a resource pool only. Virtual machines that are not currently running do not use any CPU resources. Reserved resources are not wasted if they are not used. If the utilization is less than the reservation, the resources can be utilized by other running virtual machines.
CPU available in resource pool	Indicates the total amount of resources available to satisfy a reservation for a child resource pool.	Mhz	In the under-committed state, this value is limited by the capacity at the root node. In the overcommitted case, this could be higher since we do not perform the dynamic capacity checks. Consider a resource pool with reservation=2GHz that is totally idle. It

Measurement	Description	Measurement Unit	Interpretation
			has 2GHz reserved, but it is not actually using any of its reservation. In such a case: • Other resource pools cannot reserve these 2GHz. • Other resource can use these 2GHz, that is, idle CPU reservations are not wasted.
CPU allocation used	Indicates the percentage of allocated CPU resources that are being used by this resource pool.	Percent	A high value of this measure indicates that one/more VMs in the pool are consuming the allocated CPU resources excessively.
CPU limit of pool	Indicates the limit beyond which this resource pool should not use CPU resources, even if there are resources available.	Mhz	If the *Unlimited* flag is enabled for a resource pool, then this measure will report the value *Unlimited*. This setting helps avoid wastage of idle resources. This is typically used to ensure a consistent performance of virtual machines / resource pools independent of available resources. If set to -1, then there is no fixed limit on resource usage (only bounded by available resources and shares).
Memory reserved for VMs in pool	The total amount of memory reserved by VMs in this pool.	MB	Use the detailed diagnosis of this measure to know how much memory has been reserved by each VM in the pool.
Memory consumed by VMs in pool	Indicates the total amount of memory consumed by the VMs in this pool.	MB	Use the detailed diagnosis of this measure to know how much memory has been consumed by each VM in the pool. This way, you can quickly identify memory-intensive VMs.
Memory used by running VMs	Indicates the amount of memory used by running	MB	

Measurement	Description	Measurement Unit	Interpretation
	VMs in this resource pool.		
Memory capacity	Indicates the amount of memory resources that is guaranteed available to the virtual machine or resource pool.	MB	Typically, the **Reservation** setting for a resource pool specifies the minimum acceptable amount of CPU or memory - not the amount you would like to have available. Follow the guidelines mentioned below to define the Reservation setting: • Do not set the Reservation too high. A very high reservation can limit the number of virtual machines you can power on in a resource pool. • Always leave some headroom while reserving resource. Do not commit all resources. As you move closer to fully reserving all capacity in the system, it becomes increasingly difficult to make changes to reservations and to the resource pool hierarchy without violating admission control. • If you expect frequent changes to the total available resources, use Shares, not Reservation, to allocate resources fairly across resource pools. Reserved resources are not wasted if they are not used. If the utilization is less than the reservation, the resources can be utilized by other running virtual machines.

Measurement	Description	Measurement Unit	Interpretation
Memory capacity reserved	Indicates the total amount of memory resources that have been used to satisfy the reservation requirements of all descendants of this resource pool (includes both resource pools and virtual machines).	MB	Memory resources are typically reserved for powered-on virtual machines in a resource pool only. Virtual machines that are not currently running do not use any memory resources. Reserved resources are not wasted if they are not used. If the utilization is less than the reservation, the resources can be utilized by other running virtual machines.
Memory available in resource pool	Indicates the total amount of resources available to satisfy a reservation for a child resource pool.	MB	In the undercommitted state, this value is limited by the capacity at the root node. In the overcommitted case, this could be higher since we do not perform the dynamic capacity checks. If the resource pool does not use the reserved memory, then remember the following: • Other resource pools cannot reserve these memory resources • Other resource pools can use these memory resources, that is, idle memory reservations are not wasted.
Free memory in pool	Indicates the percentage of free memory in this resource pool.	Percent	A consistent decrease in the value of this measure could indicate a potential memory contention in the pool. Compare the value of this measure across resource pools to isolate the resource pool with the least amount of free memory.
Memory limit of pool	Indicates the limit beyond which this resource pool should not use memory resources, even if there are resources available.	MB	If the Unlimited flag is enabled for a resource pool, then this measure will report the value Unlimited. This setting helps avoid wastage of idle resources.

Measurement	Description	Measurement Unit	Interpretation
			This is typically used to ensure a consistent performance of virtual machines / resource pools independent of available resources. If set to -1, then there is no fixed limit on resource usage (only bounded by available resources and shares).
Active memory	Indicates the amount of memory resources actively used from this pool.	MB	High memory consumption by a resource pool, besides eroding the physical memory resources of the host, can also impact the memory allocations to other resource pools/virtual machines; the lack of adequate memory to a VM can significantly strain its performance.
Active memory used	Indicates the percentage of total configured or available memory in this pool.	Percent	
Zero memory	Indicates the amount of memory that is zeroed out.	MB	
Swap memory	Indicates the amount of memory that is swapped.	MB	ESX Server hosts use swapping to forcibly reclaim memory from a virtual machine when no vmmemctl driver is available because the vmmemctl driver: • Was never installed • Has been explicitly disabled • Is not running (for example, while the guest operating system is booting) • Is temporarily unable to reclaim memory quickly enough to satisfy current system demands Standard demand-paging techniques swap pages back in when the virtual machine needs them. Swap space must be reserved on disk

Measurement	Description	Measurement Unit	Interpretation
			for any unreserved virtual machine memory. This swap reservation is required to ensure the system is able to preserve virtual machine memory under any circumstances. In practice, only a small fraction of the swap space may actually be used. Typically, swap space usage for each VM should be low. Since access from RAM is much faster than access from physical disk, excessive usage of swap memory will slow down the performance of a VM. Watch for VMs that are seeing higher swap usage and more swap reads and writes.
Memory overhead	Indicates the memory overhead incurred by this resource pool.	MB	ESX Server virtual machines can incur two kinds of memory overhead: • The additional time to access memory within a virtual machine. • The extra space needed by the ESX Server host for its own code and data structures, beyond the memory allocated to each virtual machine. The ESX server's memory virtualization ensures that the time overhead to memory accesses is minimal. The memory space overhead, on the other hand, is composed of two other components: • A fixed system-wide overhead for the service console (in the case of ESX 3/3.5) and the VMkernel. • Additional overhead for each virtual

Measurement	Description	Measurement Unit	Interpretation
			machine In addition, the space reserved for the virtual machine frame buffer and various virtualization data structures, also add to the memory overhead. Addition of more virtual CPUs, allocation of more memory to the guests, and choosing to configure 32-bit guest operating systems instead of 64-bit ones, can go a long way in reducing this overhead. The ESX server also provides optimizations such as memory sharing to save up memory.
Balloon memory	Indicates the amount of memory used by memory control.	MB	The vmmectl driver that is installed on a virtual machine, emulates an increase or decrease in memory pressure on the guest operating system; this way, it forces the guest OS to place memory pages into its local swap file. This driver differs from the VMware swap file method as it forces the operating system to determine what memory it wishes to page. Once the memory is paged locally on the guest operating system, the free physical pages of memory may be reallocated to other guests. As the ESX hosts sees that memory demand has been reduced, it will instruct vmmemctl to “deflate” the balloon and reduce pressure on the guest OS to page memory. The maximum amount of memory that can be reclaimed from a guest may be configured by modifying the “sched.mem.maxmemctl” advanced option. If the memory reclaimed from a guest

Measurement	Description	Measurement Unit	Interpretation
			(i.e., the value of this measure) is very low, it indicates excessive memory usage by the guest. Under such circumstances, you might want to consider allocating more memory to the guest.
Shared memory	Indicates the amount of memory in the resource pool that is shared.	MB	
Memory granted	Indicates the amount of memory that is granted to this resource pool.	MB	

3.5 The Datacenters Layer

The **Datacenters** layer monitors the datacenters managed by vCenter and reports the number of ESX servers, VMs, networks, and datastores available in each datacenter. The layer also auto-discovers the datastores in each datacenter and reports the availability and disk space usage of each datastore.

Figure 3.35: The tests mapped to the Datacenters layer

3.5.1 Datastores Test

This test auto-discovers the datastores available in each datacenter managed by VC, and reports the availability and resource usage of every datastore.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the every *datacenter:datastore* combination being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures*document. To ensure that the test extracts updated datastore metrics from the vCenter server, you need to configure the credentials of vCenter user with permissions to **Browse Datastore**.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-

Parameter	Description
	enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Report Datacenter	By default, this test reports the metrics for each *datacenter:datastore* combination. Therefore, this parameter is set to **Yes**.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Esx servers using the datastore	Indicates the number of ESX servers that are currently using this datastore.	Number	A single datastore can be used by both the ESX servers and the VMs executing on them. Use the detailed diagnosis capability of each of these measures to know which ESX servers and VMs are currently using this datastore.
VMs using the datastore	Indicates the number of VMs that are currently using this datastore.	Number	
Physical disk capacity	Indicates the total capacity of the	MB	

Measurement	Description	Measurement Unit	Interpretation
	datastore.		
Used space	Indicates the space in the datastore that is currently in use.	MB	
Free space	Indicates the space currently available for use in the datastore.	MB	
Physical disk space usage	Indicates the percentage of datastore space utilized.	Percent	A high value is indicative of excessive usage of disk space by this datastore. Comparing the value of this measure across datastores will enable administrators identify resource-intensive datastores.
Datastore availability	Indicates whether this datastore is available or not.	Percent	While the value 0 indicates that the datastore is not available, the value 100 indicates that the datastore is currently available. If a datastore becomes unavailable, then VMs that are currently using that datastore could be rendered inaccessible to users. Using the detailed diagnosis of this measure, you can identify the vSphere hosts that are using this datastore, the LUN name and drive type of the datastore, and can also ascertain how space in the datastore is being utilized.
Allocated space	Indicates the amount of physical space provisioned by an administrator for this datastore.	MB	This measure reports the storage size up to which files on this datastore can be stacked.
Average write requests rate	Indicates the average number of write commands issued per second to this datastore.	Commands/Sec	By comparing the values of these measures across datastores, you can accurately identify those datastores that are overloaded with read/write requests.
Average read	Indicates the average	Commands/Sec	

Measurement	Description	Measurement Unit	Interpretation
requests rate	number of read commands issued per second to this datastore.		
Write latency	Indicates the average amount of time taken to write to this datastore.	Secs	Very high values for these measures are indicative of processing bottlenecks in the datastore. Compare the values of these measures across datastores to determine which datastore is experiencing the maximum latency.
Read latency	Indicates the average amount of time taken to read from this datastore.	Secs	
Read rate	Indicates the rate at which data is read from this datastore.	MB/Sec	
Write rate	Indicates the rate at which data is written to this datastore.	MB/Sec	
Storage I/O control normalized latency	Indicates the device latency that the host observes while communicating with this datastore; this latency is a normalized average across virtual machines	Secs	Storage I/O Control allows cluster-wide storage I/O prioritization, which allows better workload consolidation and helps reduce extra costs associated with over provisioning. When you enable Storage I/O Control on a datastore, ESX/ESXi begins to monitor the device latency that hosts observe when communicating with that datastore. When device latency exceeds a threshold, the datastore is considered to be congested and each virtual machine that accesses that datastore is allocated I/O resources in proportion to their shares. You set shares per virtual machine. You can adjust the number for each based on need.
Storage I/O control aggregated IOPS	Indicates the aggregate number of I/O operations that are occurring on this datastore.	Number	

Measurement	Description	Measurement Unit	Interpretation
Storage I/O Control active time	Indicates whether storage I/O control (SIOC) is active on this datastore or not.	Percent	Storage I/O Control (SIOC) is used to provide I/O prioritization of virtual machines running on a group of VMware vSphere hosts that have access to a shared storage pool. It extends the familiar constructs of shares and limits, which exist for CPU and memory, to address storage utilization through a dynamic allocation of I/O capacity across a cluster of vSphere hosts. If SIOC is active, the value of this measure will be 100%. If it is not active, the value of this measure will be 0.
Total IOPS	Indicates the total number of read and write commands issued per second to this datastore.	Commands/Sec	Compare the value of this measure across datastores to identify the busiest datastore in terms of the rate at which read/write commands are issued on it. A consistent increase in the value of this measure for a particular datastore could indicate a potential I/O overload.
Total latency	Indicates the total amount of time taken to read and write to this datastore.	Secs	Ideally, the value of this measure should be low. A high value could indicate an I/O processing bottleneck. By comparing the value of this measure across datastore, you can quickly identify which datastore is the most latent. Such datastores are ideal candidates for enabling SIOC. EnablingIOC on a datastore triggers the the monitoring of device latency that hosts observe when communicating with that datastore. When latency exceeds a set threshold the feature engages automatically as the datastore is experiencing congestion. Each virtual machine that accesses that datastore is then allocated I/O resources in proportion to their shares.

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Throughput</td><td>Indicates the rate at which the data is read and written to this datastore.</td><td>MB/Sec</td><td>Ideally, the value of this measure should be high. A steady decrease in the value of this measure could indicate that the datastore is experiencing a congestion. You may want to enable SIOC on such datastores, so that when latency exceeds a set threshold the SIOC feature engages automatically and then allocates I/O resources to each VM that acceses that datastore in proportion to their shares.</td></tr>
<tr><td>Type</td><td>Indicates the type of file system this datastore supports.</td><td></td><td>The values that this measure can report are listed below, along with their descriptions:<table><tr><th>Measure Value</th><th>Description</th></tr><tr><td>VMFS</td><td>VMFS (Virtual Machine File system) is a high-performance cluster file system optimized for virtual machines. While conventional file systems allow only one server to have read-write access to the same file system at a given time, VMFS leverages shared storage to allow multiple VMware vSphere hosts to read and write to the same storage concurrently.</td></tr></table></td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation

Measure Value	Description
NFS	This is a file system on a NAS storage device. The vSphere host can access a designated NFS volume located on an NFS server, mount the volume, and use it for any storage needs.
NFSV41	NFS v4.1 introduces better performance and availability through load balancing and multipathing. Another major enhancement with NFS v4.1 is the security aspect. With this version, Kerberos and thus non-root user authentication are both supported.
CIFS	CIFS is the standard way that computer users share files across corporate intranets and the Internet.
VFAT	Virtual File Allocation Table (VFAT) handles long file

Measurement	Description	Measurement Unit	Interpretation	
			Measure Value	**Description**
				names, which otherwise could not be handled by the original file allocation table (FAT) programming. A file allocation table is the means by which the operating system keeps track of where the pieces of a file are stored on a hard disk. Since the original FAT assumed file names were limited to a length of eight characters, a program extension was needed to handle the longer names. Microsoft refers to this extension as a driver. The VFAT extension runs in protected mode, uses 32-bit code, and uses VCACHE for disk cache.
			VSAN	Virtual SAN virtualizes local physical storage resources of vSphere hosts and turns them

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Description</th></tr>
<tr><td></td><td>into pools of storage that can be carved up and assigned to virtual machines and applications according to their quality of service requirements.</td></tr>
<tr><td>VFFS</td><td>VFFS (Virtual Flash File System) is a derivative of VMFS, which is optimized for SSDs and is used to group the physical SSDs into a single caching resource pool. It is a non-persistent resource and, therefore, virtual machines cannot be stored in it. A VFFS resource pool can be created from mixed resources. All devices types are treated the same and no distinction is made between SAS, SATA, PCI express connectivity.</td></tr>
</table>

The numeric values that correspond to each of the Measure Values listed above are as follows:
</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>VMFS</td><td>1</td></tr><tr><td>NFS</td><td>2</td></tr><tr><td>NFSV41</td><td>3</td></tr><tr><td>CIFS</td><td>4</td></tr><tr><td>VFAT</td><td>5</td></tr><tr><td>VSAN</td><td>6</td></tr><tr><td>VFFS</td><td>7</td></tr></table> **Note:** By default, this measure reports the **Measure Value**s listed in the table above. In the graph of this measure however, the file system type of the datastore is represented using the corresponding numeric equivalents only.

The detailed diagnosis of the *ESX servers using the datastore* measure lists the ESX servers that are using this datastore.

Details of ESX servers present in the datastore

Time	VMName	EsxHostName	EsxHostIP
Mar 30, 2009 10:20:05	-	esx3i	192.168.10.179

Figure 3.36: The detailed diagnosis of the ESX servers using the datastore measure

The detailed diagnosis of the *VMs using the datastore* measure lists the VMs that are currently using this datastore.

Details of Virtual machines present in the datastore

Time	VMName	EsxHostName	EsxHostIP
Mar 30, 2009 10:20:05			
	xp3 logmein and gotomypc	esx3i	192.168.10.179
	xp2(oracle10.60)	esx3i	192.168.10.179
	leostream connection broker	esx3i	192.168.10.179
	leostream cb	esx3i	192.168.10.179

Figure 3.37: The detailed diagnosis of the VMs using the datastore measure

3.5.2 Datacenters Test

A **datacenter** is the primary container of inventory objects such as hosts and virtual machines. From the datacenter, you can add and organize inventory objects. Typically, you can add hosts and folders, to a datacenter.

This test auto-discovers the datacenters managed by VC and reports on the configuration of each datacenter – i.e., how many ESX hosts, VMs, networks, clusters, and datastores are available in every datacenter.

Target of the test : A VMware vCenter server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the every *datacenter:datastore* combination being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The host for which this test is to be configured.
Port	Refers to the port at which the specified host listens to.
VC User and VC Password	To connect to vCenter and extract metrics from it, this test should be configured with the name and password of a user with **Administrator** or **Virtual Machine Administrator** privileges to vCenter. However, if, owing to security constraints, you are not able to use the credentials of such users for test configuration, then you can configure this test with the credentials of a user with **Read-only** rights to vCenter. For this purpose, you can assign the '**Read-only**' role to a local/domain user to vCenter, and then specify name and password of this user against the VC User and VC Password text boxes. The steps for assigning this role to a user on vCenter have been detailed in the *Monitoring VMware Infrastructures* document. vCenter servers terminate user sessions based on timeout periods. The default timeout period is 30 mins. When you stop an agent, sessions currently in use by the agent will remain open for this timeout period until vCenter times out the session. If the agent is restarted within the timeout period, it will open a new set of sessions. If you want the eG agent to close already existing sessions on vCenter before it opens new sessions, then, instead of the '**Read-only**' user, you can optionally configure the VC User and VC Password parameters with the credentials of a user with permissions to View and Stop

Parameter	Description
	Sessions on vCenter. For this purpose, you can create a special role on vCenter, grant the View and Stop Sessions privilege (prior to vCenter 4.1, this was called the View and Terminate Sessions privilege) to this role, and then assign the new role to a local/domain user to vCenter. The steps for this have been discussed in the *Monitoring VMware Infrastructures* document.
Confirm Password	Confirm the password by retyping it in this text box.
SSL	By default, the vCenter server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the vCenter server via HTTPS by default.
Webport	By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while monitoring vCenter, the eG agent, by default, connects to port 80 or 443, depending upon the SSL-enabled status of vCenter – i.e., if vCenter is not SSL-enabled (i.e., if the SSL flag above is set to **No**), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled (i.e., if the ssl flag is set to **Yes**), then the agent-vCenter communication occurs via port 443 by default. Accordingly, the Webport parameter is set to **default** by default. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the Webport parameter, you can specify the exact port at which vCenter in your environment listens, so that the eG agent communicates with that port for collecting metrics from vCenter.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
ESX servers in datacenter	Indicates the number of ESX servers in this datacenter.	Number	
ESX servers directly assigned to datacenter	Indicates the number of ESX servers that are directly connected to the datacenter.	Number	The detailed diagnosis of this measure, if enabled, lists those ESX servers that are directly assigned to this datacenter – i.e., those that are not under clusters.
ESX servers under clusters in datacenter	Indicates the number of ESX servers in the datacenter that are under clusters.	Number	The detailed diagnosis of this measure, if enabled, lists the ESX servers that are under clusters.
VMs in datacenter	Indicates the number of VMs in the datacenter.	Number	
VMs directly assigned to servers in datacenter	Indicates the number of VMs directly assigned to this datacenter.	Number	The detailed diagnosis of this measure, if enabled, lists the VMs that directly assigned to this datacenter – i.e., those that are not under a cluster.
VMs under clusters in datacenter	Indicates the number of VMs in a cluster.	Number	The detailed diagnosis of this measure, if enabled, lists the VMs that are in a cluster.
Clusters in datacenter	Indicates the number of clusters configured within a datacenter.	Number	The detailed diagnosis of this measure, if enabled, lists the clusters in the datacenter, and the IP address and name of the ESX servers that are under each cluster.
Networks in datacenter	Indicates the number of networks in this datacenter.	Number	The detailed diagnosis of this measure, if enabled, lists the networks in the datacenter.
Datastores in datacenter	Indicates the number of datastores in this datacenter.	Number	The detailed diagnosis of this measure, if enabled, lists the datastores in the datacenter.

The detailed diagnosis of the *Esx servers under cluster* measure lists the name and IP address of ESX servers that are under clusters configured on datacenters.

Details of ESX servers under the cluster in Datacenter			
Time	GuestName	EsxHostName	EsxHostIP
Mar 27, 2009 12:12:53			
	-	esx3	192.168.10.136
	-	esx3i	192.168.10.179

Figure 3.38: The detailed diagnosis of the Esx servers under cluster

The detailed diagnosis of the *Directly connected vms* measure lists the name of the guests that are directly assigned to the datacenter, IP address of ESX server that hosts the guest, and the name of the ESX server.

Details of directly connected VMs in Datacenter			
Time	GuestName	EsxHostName	EsxHostIP
Mar 27, 2009 12:12:53			
	esxreadytesttemplate	esx3	192.168.10.136
	new template	esx3	192.168.10.136
	redhat _sherief	esx3	192.168.10.136
	redhat_ linux _sherief	esx3	192.168.10.136
	vmware studio	esx3	192.168.10.136
	eg vm ready appliance	esx3	192.168.10.136
	jk-temp	esx3	192.168.10.136
	virtual center 2003	esx3	192.168.10.136
	winsys-karthi	esx3	192.168.10.136

Figure 3.39: The detailed diagnosis of the Directly connected vms measure

The detailed diagnosis of the *VMs under cluster* measure lists the name of the guests under a cluster, IP address of ESX server that hosts the guest, and the name of the ESX server.

Details of Vms under the cluster in Datacenter			
Time	GuestName	EsxHostName	EsxHostIP
Mar 27, 2009 12:12:53			
	wind2003eur	esx3	192.168.10.136
	esxreadytesttemplate	esx3	192.168.10.136
	redhat_ linux _sherief	esx3	192.168.10.136
	winsys-karthi	esx3	192.168.10.136
	eg vm ready appliance	esx3	192.168.10.136
	redhat _sherief	esx3	192.168.10.136
	dummy for soban	esx3	192.168.10.136
	virtual center 2003	esx3	192.168.10.136
	vmware studio	esx3	192.168.10.136
	leostream connection broker	esx3i	192.168.10.179
	linel4-dhanabal	esx3i	192.168.10.179
	xp3 logmein and gotomypc	esx3i	192.168.10.179
	xp2(oracle10.60)	esx3i	192.168.10.179
	leostream cb	esx3i	192.168.10.179

Figure 3.40: The detailed diagnosis of the VMs under cluster measure

Chapter 4: Monitoring vCenter Clusters

Typically, when one/more vCenter servers in an environment are grouped to form a cluster, all requests to the vCenter servers in the cluster will be routed via a vCenter cluster server. If one vCenter server in the cluster is unavailable, then the cluster server will automatically direct the incoming request to any other available vCenter server in the cluster; this is why, all the vCenter servers in the cluster are required to manage the same set of vSphere/ESX servers. This way, the request does not go unserviced and the 100% availability of the target vCenter server is ensured.

However, if none of the vCenter servers in a cluster are available, then the cluster server will not be able to service any of the requests.

To enable administrators to determine whether the cluster is able to service requests or not, the eG Enterprise Suite includes the VMware vCenter Cluster component.

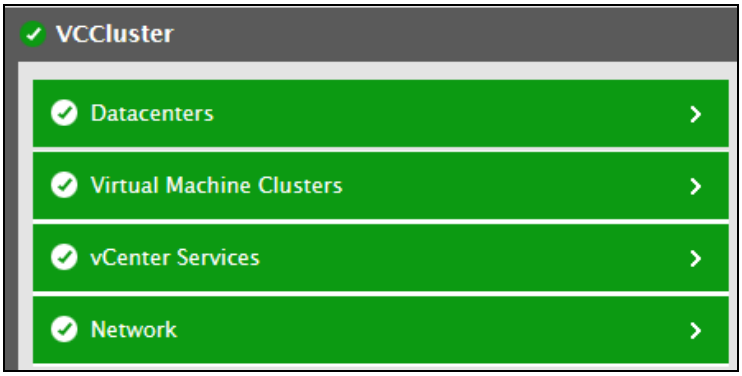

Figure 4.1: The VMware vCenter Cluster monitoring model

When the eG agent monitoring this component executes tests on it, the cluster server automatically connects to any available vCenter server in the cluster and pulls out the desired metrics. This is why the layer model of the cluster server is the same as that of a vCenter server.

If no server is currently available, then none of the tests will work, thus indicating that the cluster is not alive.

About eG Innovations

eG Innovations provides intelligent performance management solutions that automate and dramatically accelerate the discovery, diagnosis, and resolution of IT performance issues in on-premises, cloud and hybrid environments. Where traditional monitoring tools often fail to provide insight into the performance drivers of business services and user experience, eG Innovations provides total performance visibility across every layer and every tier of the IT infrastructure that supports the business service chain. From desktops to applications, from servers to network and storage, from virtualization to cloud, eG Innovations helps companies proactively discover, instantly diagnose, and rapidly resolve even the most challenging performance and user experience issues.

eG Innovations is dedicated to helping businesses across the globe transform IT service delivery into a competitive advantage and a center for productivity, growth and profit. Many of the world's largest businesses use eG Enterprise to enhance IT service performance, increase operational efficiency, ensure IT effectiveness and deliver on the ROI promise of transformational IT investments across physical, virtual and cloud environments.

To learn more visit www.eginnovations.com.

Contact Us

For support queries, email support@eginnovations.com.

To contact eG Innovations sales team, email sales@eginnovations.com.