# Monitoring VMware vCenter

## eG Enterprise v6

# Table of Contents

Chapter

**1**

# Monitoring VMware vCenter

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 depicts the architecture of vCenter.
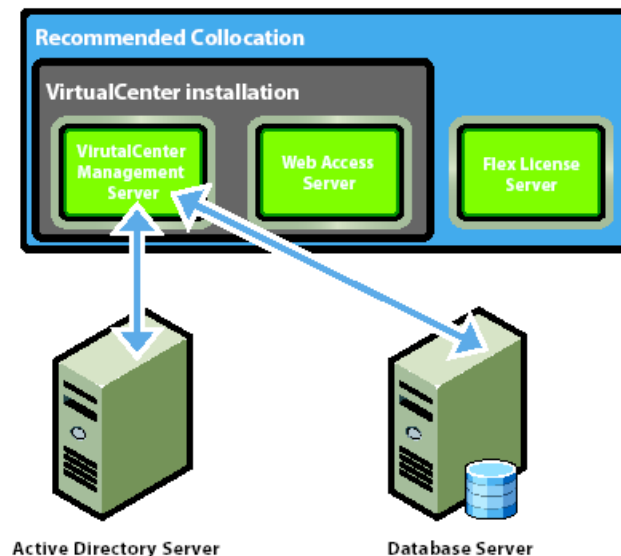


Figure 1: The architecture of vCenter

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.

eG Enterprise prescribes a specialized *VMware vCenter* monitoring model (see Figure 2), 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.
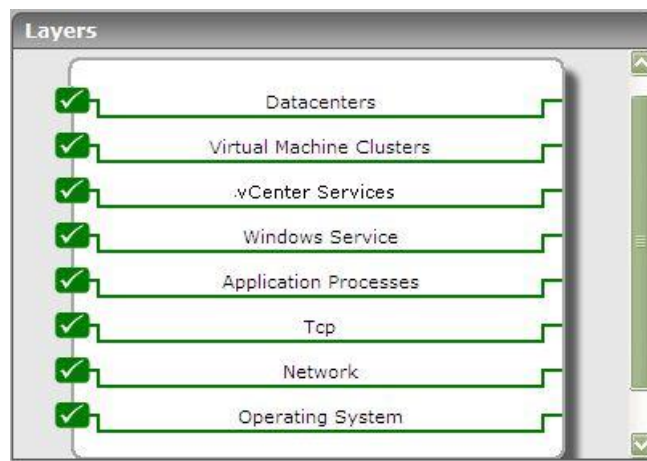


Figure 2: The layer model of vCenter

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 5 layers of Figure 2 have already been dealt with in the *Monitoring Unix and Windows Servers* document, the sections to come will deal with the top 2 layers only.

## 1.1 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
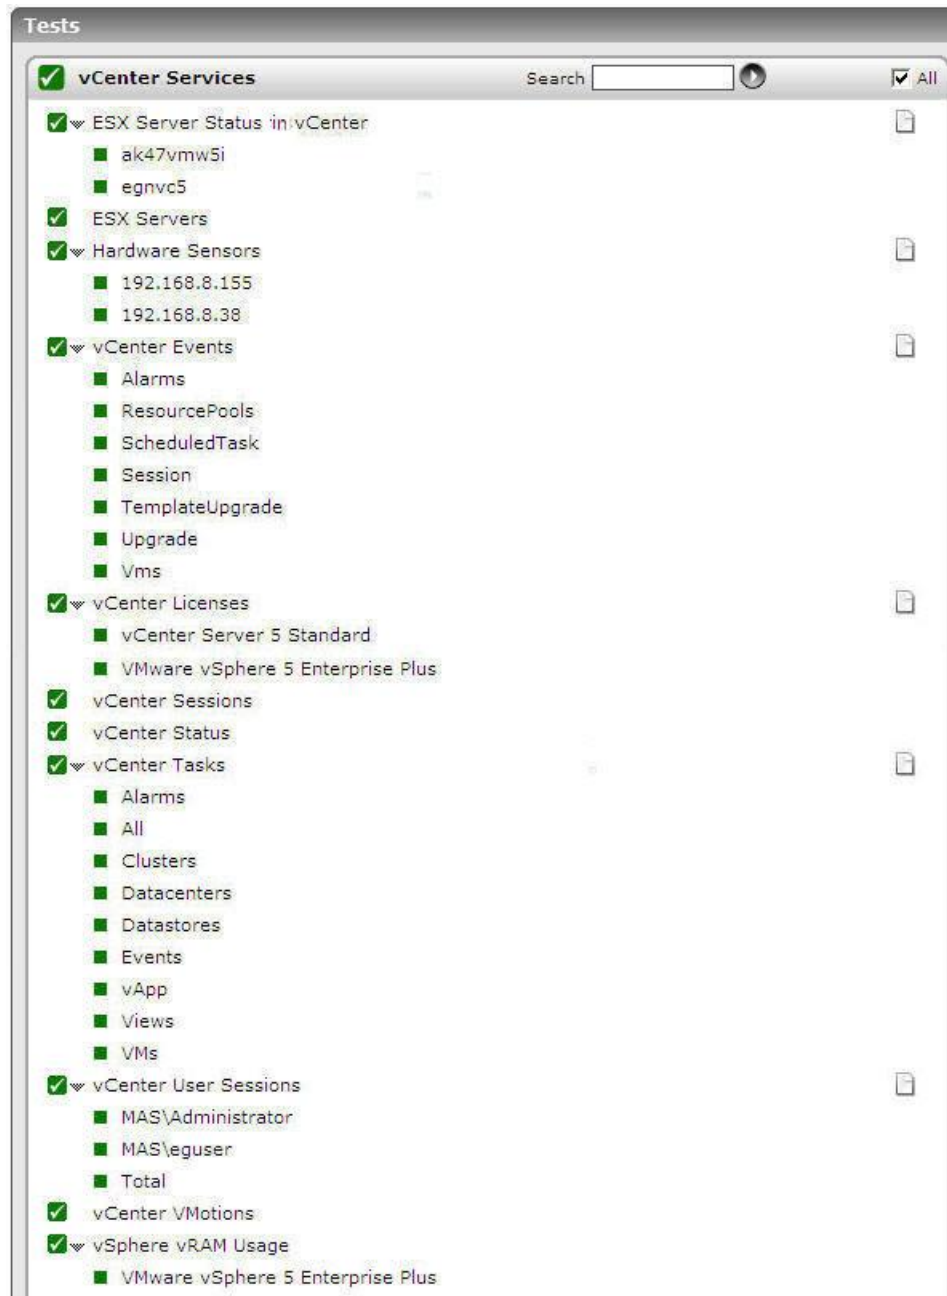
Figure 3: The tests mapped to the vCenter Services layer

## 1.1.1    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.

| Purpose | 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 | vCenter |

| Agent deploying the test | An internal agent |
|---|---|
| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed<br><br>2. **HOST** - The host for which the test is to be configured<br><br>3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL.<br><br>4. **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 Section 8.3 of the *Monitoring VMware Infrastructures* document.<br><br>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 Section 8.4 of the *Monitoring VMware Infrastructures* document.<br><br>5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box.<br><br>6. **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.<br><br>7. **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.<br><br>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. |
| Outputs of the test | One set of results for each vSphere/ESX host managed by the vCenter server being monitored |

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|---|---|---|---|
| | **Connection status:**<br><br>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:<br><br>| Measure Value | Numeric Value |<br>|---|---|<br>| Disconnected | 0 |<br>| Connected | 1 |<br>| NotResponding | 2 |<br><br>**Note:**<br><br>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:**<br><br>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.<br><br>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 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.<br><br>The master FDM is determined through election by the FDMs that are alive at the time. An election occurs in the following circumstances:<br><br>a. When the vSphere HA feature is enabled for the cluster.<br><br>b. When the master's host fails.<br><br>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.<br><br>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.<br><br>The slaves are responsible for reporting state updates to the master and restarting VMs as required.<br><br>The values that this measure reports and their numeric equivalents have been discussed in the table below |

| Measure Value | Numeric Value |
|---|---|
| Master | 0 |
| ConnectedToMaster | 1 |
| Election | 2 |
| Uninitialized | 3 |
| NetworkPartitionedFromMaster | 4 |
| FdmUnreachable | 5 |
| HostDown | 6 |
| NetworkIsolated | 7 |
| InitializationError | 8 |
| UnitializationError | 9 |

**Note:**

Typically, this measure reports one of the **Measure Value**s 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.

| **Is in maintenance mode?**<br><br>Indicates whether this host is currently in the maintenance mode or not. | | 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.<br><br>The values this measure reports and their numeric equivalents are provided in the table below: |

| **Measure Value** | **Numeric Value** |
|---|---|
| Yes | 1 |
| No | 0 |

**Note:**

Typically, this measure reports one of the **Measure Value**s 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.

| | | | |
|---|---|---|---|
| | **Standby mode:**<br><br>Indicates whether this host is currently in the standby mode or not. | | 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.<br><br>The values this measure reports and their numeric equivalents are provided in the table below: |

| Measure Value | Numeric Value |
|---|---|
| None | 0 |
| In | 1 |
| Entering | 2 |
| Exiting | 3 |

**Note**:

Typically, this measure reports one of the **Measure Value**s 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.

| | **Power state:**<br><br>Indicates whether this host is currently in the standby, powered off, or powered on mode. | | The values this measure reports and their numeric equivalents are provided in the table below: |
|---|---|---|---|

| Measure Value | Numeric Value |
|---|---|
| Unknown | 0 |
| PoweredOn | 1 |
| PoweredOff | 2 |
| StandBy | 3 |

**Note:**

Typically, this measure reports one of the **Measure Value**s 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.

| | **Availability status:**<br><br>Indicates whether this host is currently in the maintenance, standby, powered on, or powered off modes. | | The values this measure reports and their numeric equivalents are provided in the table below: |
|---|---|---|---|

| Measure Value | Numeric Value |
|---|---|
| In maintenance | 0 |
| PoweredOn | 1 |
| Unknown | 2 |
| PoweredOff | 3 |
| StandBy | 4 |

**Note:**

Typically, this measure reports one of the **Measure Value**s 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.

| | vSAN service enabled? <br><br> Indicates whether/not the vSAN service is currently enabled for this host. | | 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. <br><br> 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. <br><br> 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 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. <br><br> 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: <br><br> <table><tr><td>**Measure Value**</td><td>**Numeric Value**</td></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table> <br><br> **Note**: <br><br> 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 vSAN is represented using the corresponding numeric equivalents only. |
|---|---|---|---|

## 1.1.2    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.

| Purpose | 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 | vCenter |
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed<br><br>2. **HOST** - The host for which the test is to be configured<br><br>3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL.<br><br>4. **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 Section 8.3 of the *Monitoring VMware Infrastructures* document.<br><br>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 Section 8.4 of the *Monitoring VMware Infrastructures* document.<br><br>5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box.<br><br>6. **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. |

7.  **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.

8.  **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.

| Outputs of the test | One set of results for each vSphere/ESX host managed by the vCenter server being monitored | | |
|---|---|---|---|
| **Measurements made by the test** | **Measurement** | **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:**<br><br>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:**<br><br>Indicates the number of sensors in this vSphere/ESX host with errors. | Number | Use the detailed diagnosis of this measure to know which sensors are currently experiencing errors. |
| | **Percentage of error sensors:**<br><br>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.

| Warning sensor details | | | | |
|---|---|---|---|---|
| **Time** | **Type** | **Name** | **Details** | |
| Mar 21, 2011 17:59:22 | | | | |
| | Slot/Connector | Memory Device 5 DIMM B1 0 | Fault Status - Assert | |
| | Memory | Memory | - | |

Figure 4: The detailed diagnosis of the Warning sensors measure

## 1.1.3 vCenter Status Test

This test monitors the availability and response time of vCenter.

| Purpose | Monitors the availability and response time of vCenter |
|---|---|
| **Target of the test** | vCenter |
| **Agent deploying the test** | An internal agent |

| | |
|---|---|
| **Configurable parameters for the test** | 1.   **TEST PERIOD** - How often should the test be executed<br><br>2.   **HOST** - The host for which the test is to be configured<br><br>3.   **PORT –** Refers to the port at which vCenter listens. By default, this is NULL.<br><br>4.   **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 Section 8.3 of the *Monitoring VMware Infrastructures* document.<br><br>    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 Section 8.4 of the *Monitoring VMware Infrastructures* document.<br><br>5.   **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box.<br><br>6.   **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.<br><br>7.   **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.<br><br>    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. |
| **Outputs of the test** | One set of results for the vCenter server being monitored |

| **Measurements made by the** | **Measurement** | **Measurement Unit** | **Interpretation** |
|---|---|---|---|

| test | **Availability:**<br><br>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:**<br><br>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. |

## 1.1.4    vCenter Sessions Test

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

| Purpose | Monitors the user sessions to the vCenter, and reports new logins and sessions logging out |
|---|---|
| Target of the test | A vCenter server |
| Agent deploying the test | An internal agent |

| Configurable parameters for the test | 1. | **TEST PERIOD** - How often should the test be executed |
|---|---|---|
| | 2. | **HOST** - The host for which the test is to be configured |
| | 3. | **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 4. | **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 Section 8.3 of 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 Section 8.4 of the *Monitoring VMware Infrastructures* document. |
| | 5. | **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box. |
| | 6. | **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. |
| | 7. | **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. |

| | 8. **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.<br><br>The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:<br><br>&#10142; The eG manager license should allow the detailed diagnosis capability<br>&#10142; Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. |
|---|---|
| **Outputs of the test** | One set of results for the vCenter server being monitored |

| **Measurements made by the test** | **Measurement** | **Measurement Unit** | **Interpretation** |
|---|---|---|---|
| | **Current sessions to VC:**<br><br>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:**<br><br>Indicates the number of new logins to vCenter. | Number | |
| | **Percent new logins:**<br><br>Indicates the percentage of current sessions that logged in during the last measurement period. | Percent | |
| | **Sessions logging out:**<br><br>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.
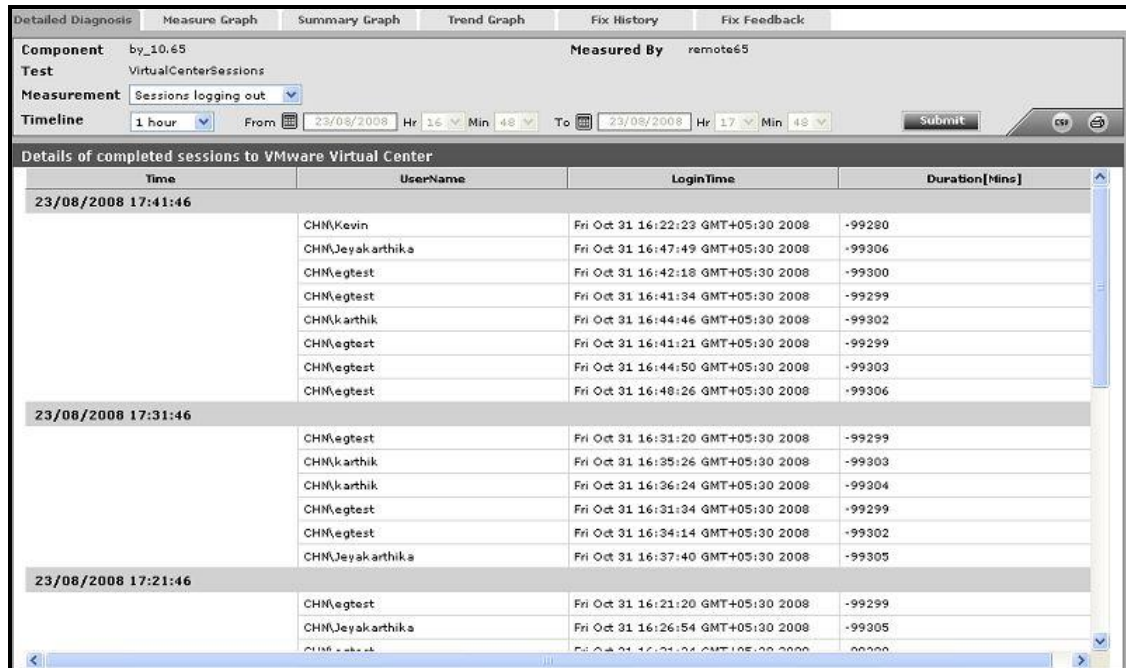
Figure 5: The detailed diagnosis of the Sessions logging out measure

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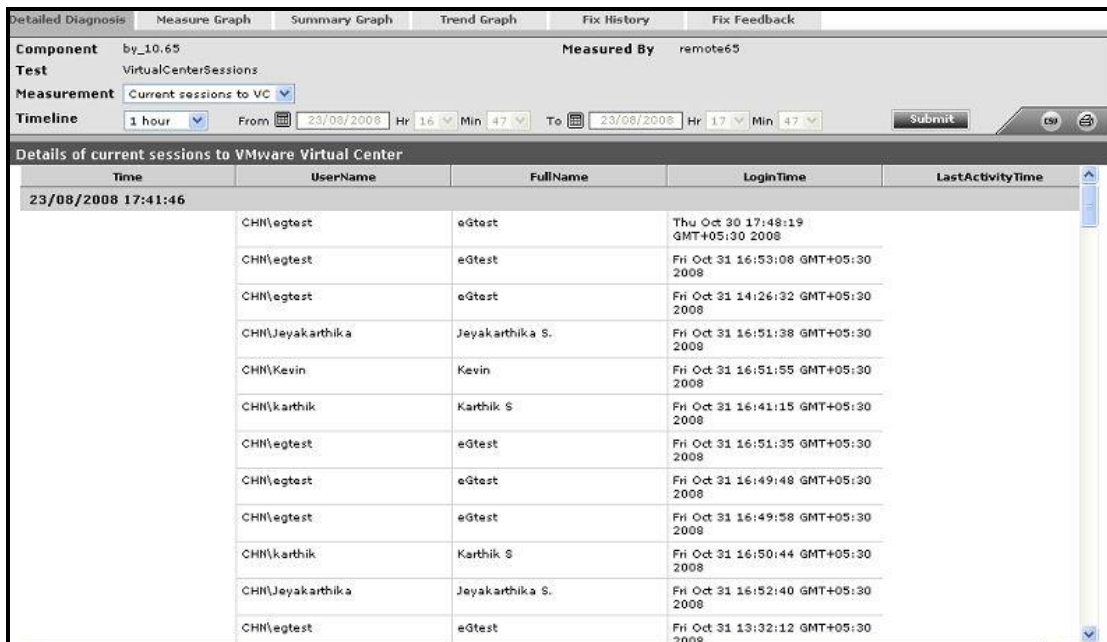


Figure 6: The detailed diagnosis of the Current sessions measure

## 1.1.5    vCenter Licenses Test

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

| Purpose | Reports critical statistics related to the license management by vCenter |
|---------|--------------------------------------------------------------------------|

| Target of the test | A vCenter server |
|---|---|
| Agent deploying the test | An internal agent |

| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed |
|---|---|
| | 2. **HOST** - The host for which the test is to be configured |
| | 3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 4. **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 *n*one. 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 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:** <br><br> vCenter 4.0 does not support the command line interface; therefore, set the **USING** parameter to **Web services API** while monitoring vCenter 4.0. |
| | 5. **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.* |
| | 6. **CONFIRM PASSWORD** – Confirm the password by retyping it here. |
| | 7. **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*. |

|  | |
|---|---|
| | 8. **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**.<br><br>9. **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.<br><br>The status of the **REREADLICENSE** flag will be disregarded if the **USING** parameter is set to **Command Line Interface**.<br><br>10. **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.<br><br>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.<br><br>11. **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.<br><br>The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:<br><br> ▪ The eG manager license should allow the detailed diagnosis capability<br><br> ▪ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. |
| **Outputs of the test** | One set of results for the vCenter server being monitored |
| **Measurements made by the test** | <table><tr><th>Measurement</th><th>Measurement Unit</th><th>Interpretation</th></tr><tr><td>**Licenses installed:**<br><br>Indicates the number of licenses currently installed on vCenter.</td><td>Number</td><td></td></tr></table> |

| | 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 of licenses currently in use. | Percent | Ideally, this value should be low. A value equal to or close to 100% is a cause for concern, as it indicates that the vCenter has run out of licenses. |

## 1.1.5.1 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 7 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 7.

Figure 7: Connecting to vCenter

4.    Figure 7 then appears. Click on the **Administration** icon indicated by Figure 8.
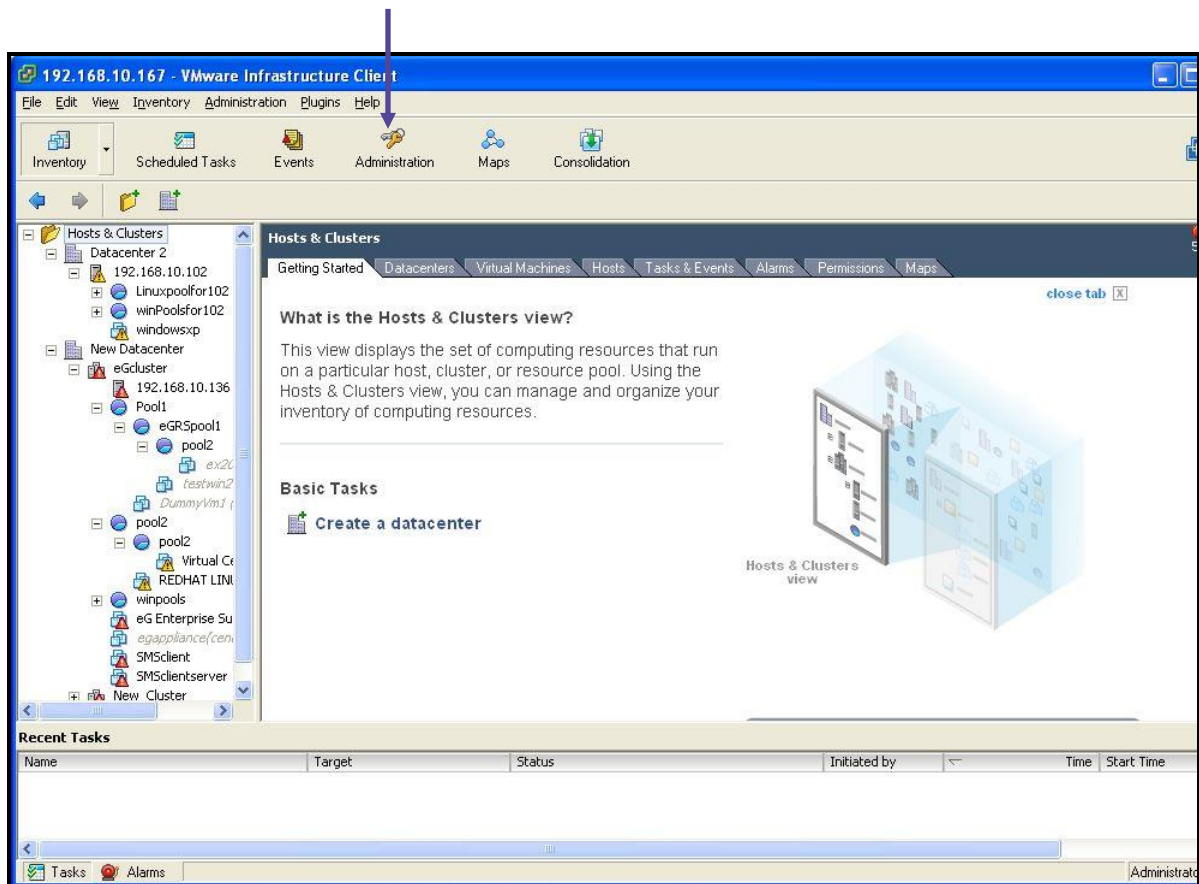


Figure 8: The VMware Infrastructure Client console

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 9.
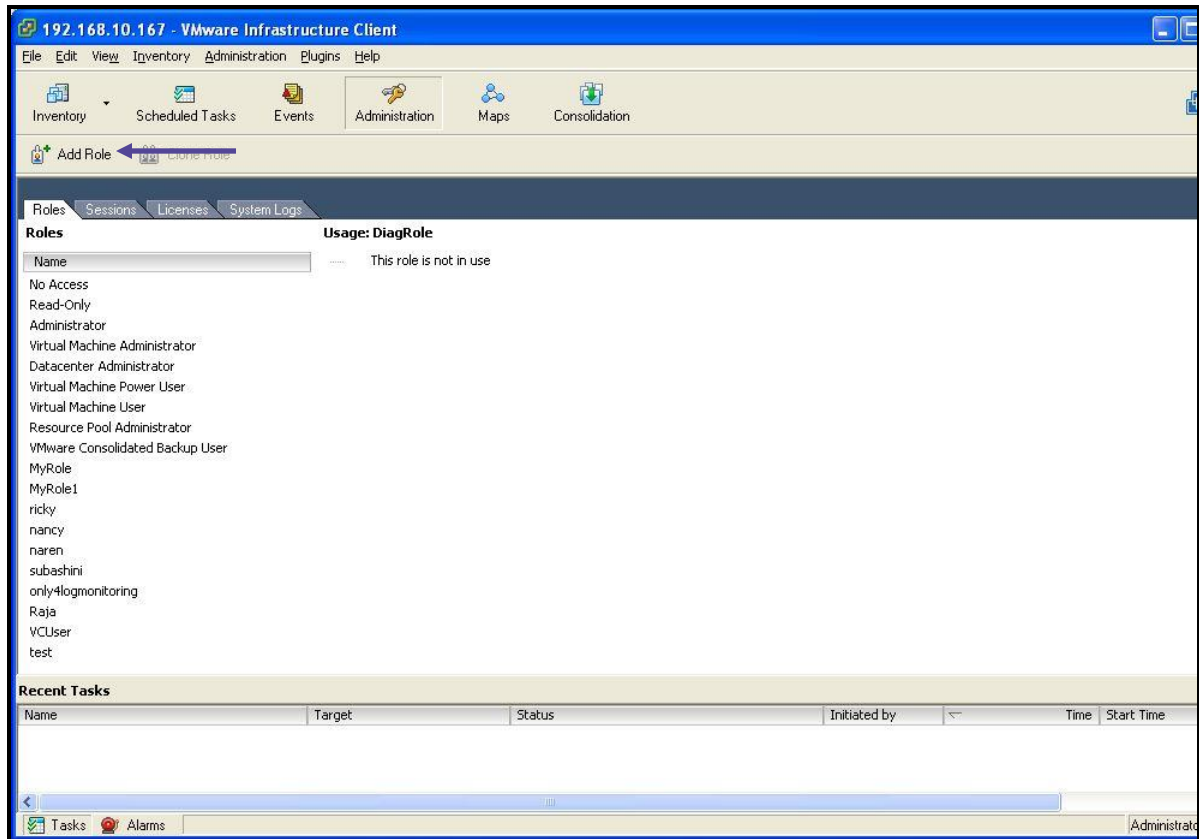
Figure 9: Adding a new role

6.   In Figure 10 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 10. To grant the **Licenses** permission to the new role, select the check box corresponding to **Licenses** in Figure 10. Then, click the **OK** button.
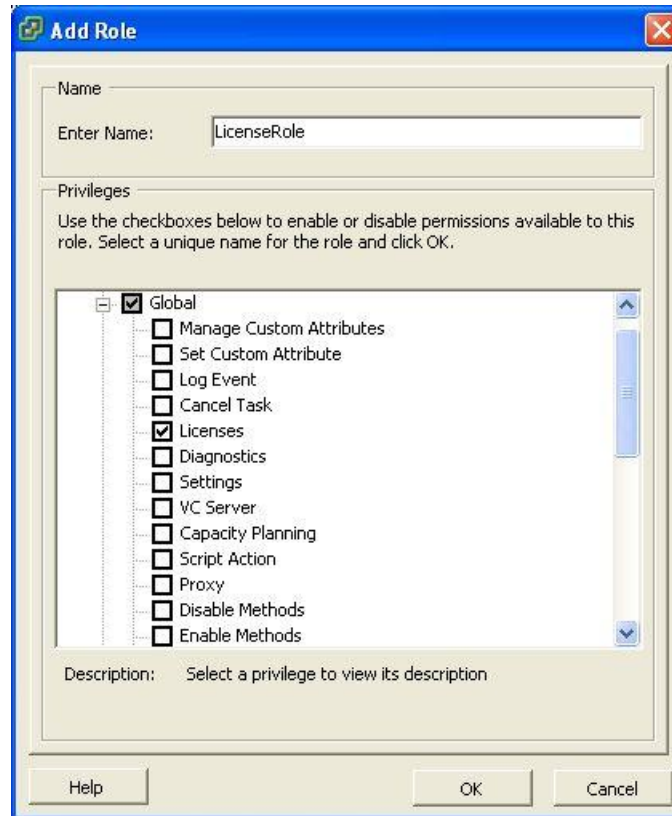
Figure 10: Providing the details of the new role

7.    Figure 11 will then appear displaying the role that you just created.
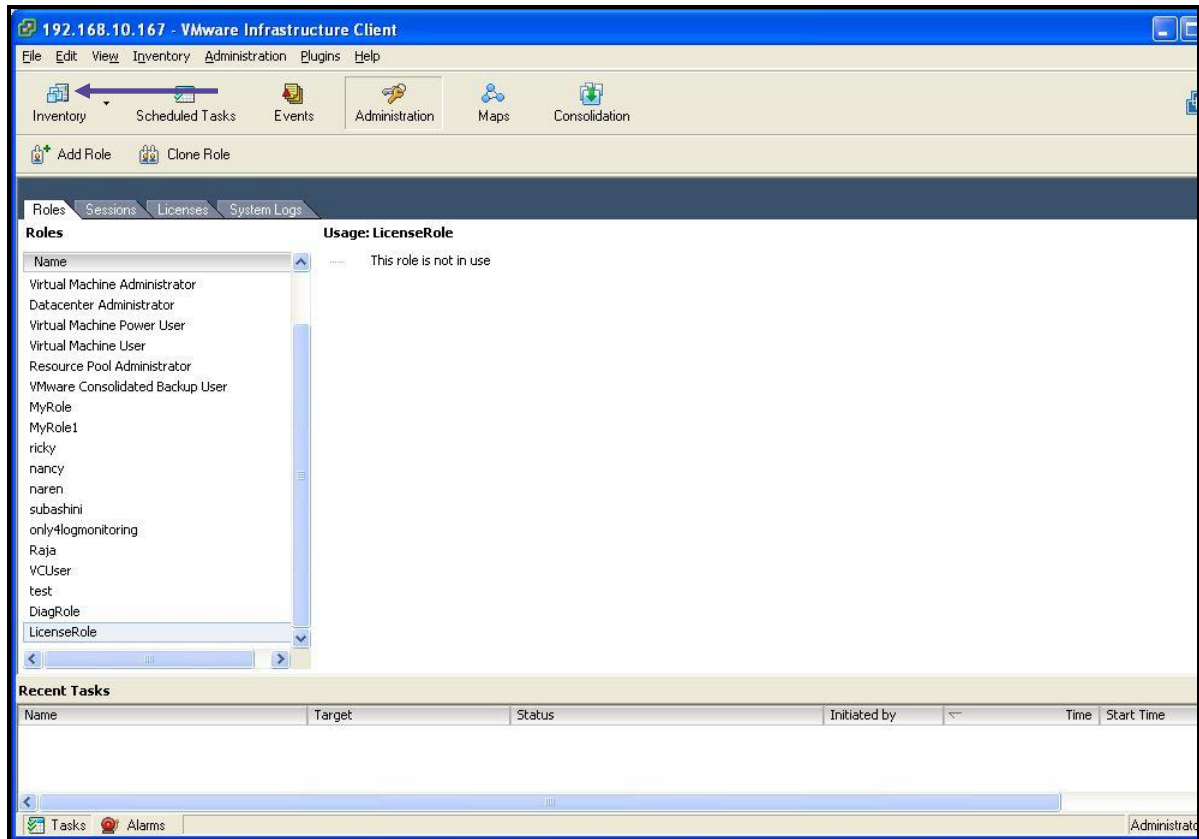
Figure 11: The newly created role appended to the list of existing roles in the Roles tab page

8.    Next, click on the **Inventory** icon indicated by Figure 11. When Figure 12 appears, click on the **Hosts & Clusters** node in the tree-structure in the left panel of Figure 12, 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.
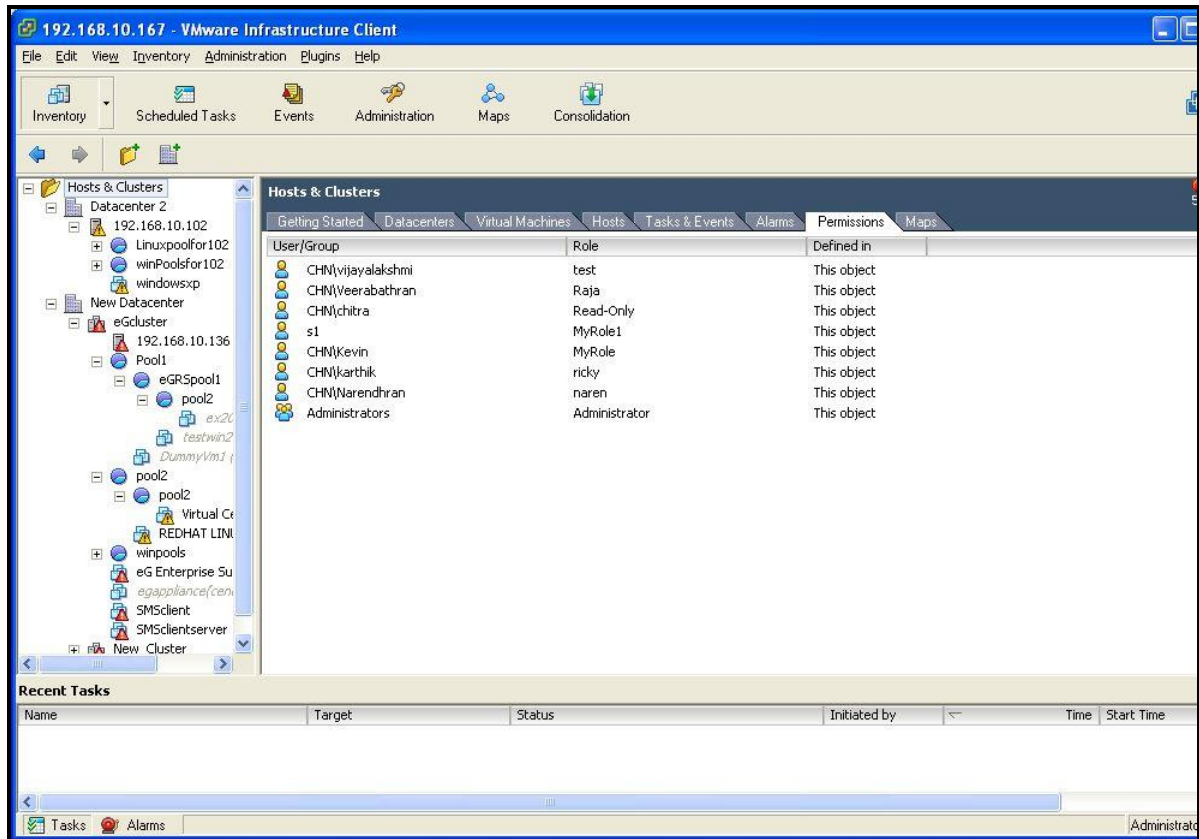
Figure 12: The Permissions tab page

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.
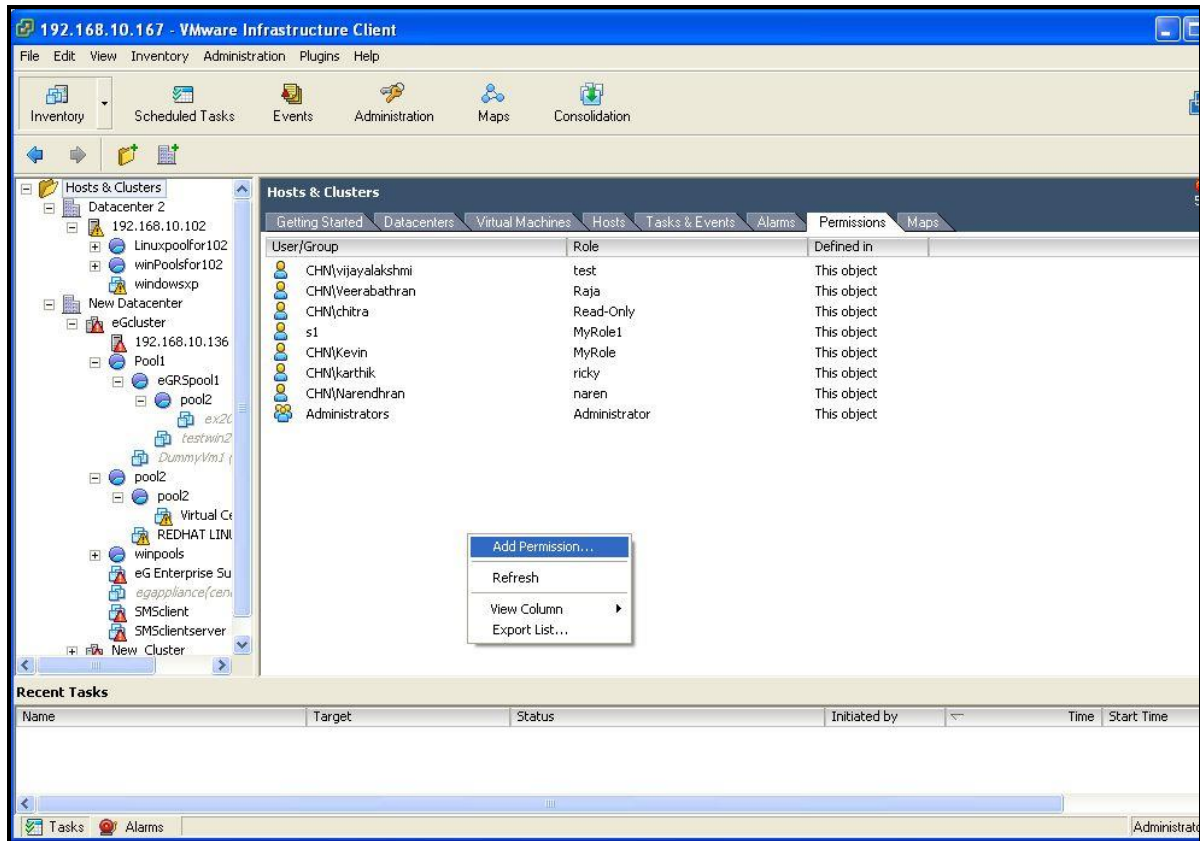
Figure 13: Selecting the Add Permission option

10. Figure 14 then appears. To choose the user who is to be assigned the new role, click the **Add** button in Figure 14.
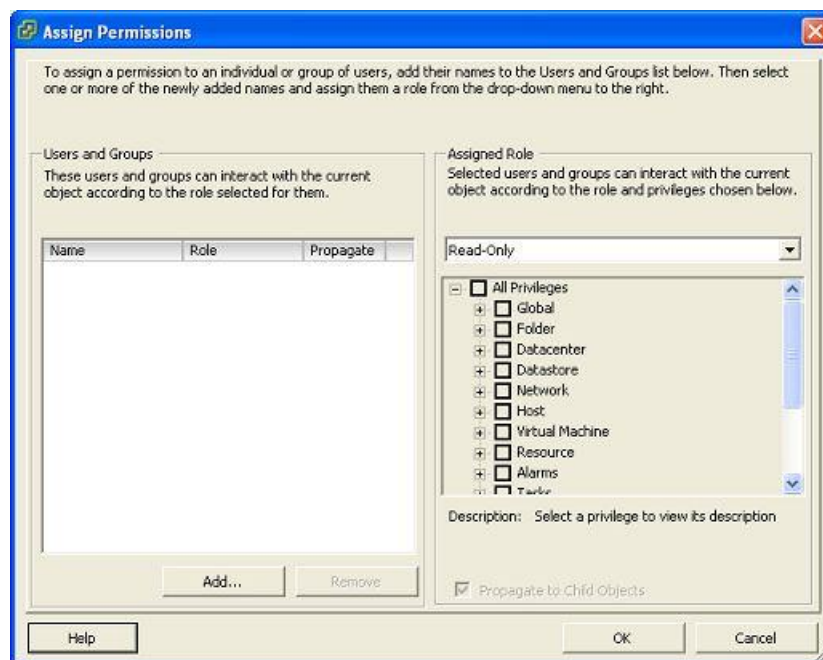


Figure 14: Adding a user

11.  This will open Figure 15. 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.

Either way, to add the user, click on the **Add** button in Figure 15. This will add the chosen user to the **Users** box. Then, click the **OK** button in Figure 15.
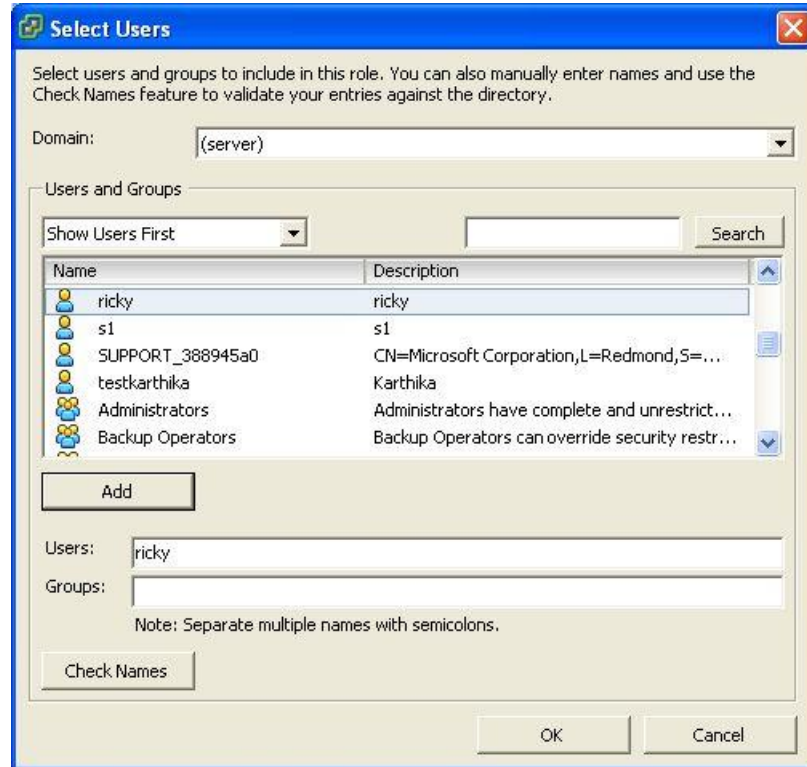
Figure 15: Selecting a local/domain user

12.  When you return to Figure 14, you will find that the local/domain user you selected from Figure 15 is displayed in the **Users and Groups** section (see Figure 16).
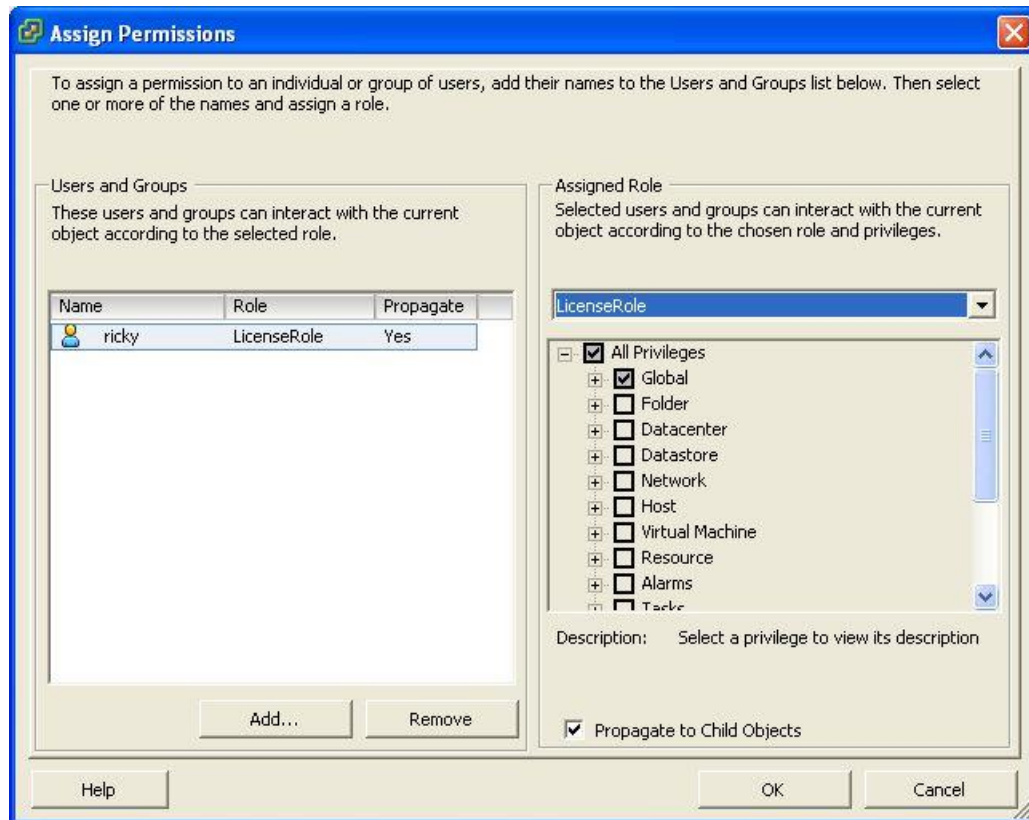
Figure 16: Assign the newly created role to the user

13. To assign the newly created role to this user, select the user displayed in the **Users and Groups** section in Figure 16, pick the newly created role from the **Assigned Role** section, and then, click the **OK** button in Figure 16.

14. Figure 17 then appears displaying the user-role mapping

MONITORING VMWARE VCENTER
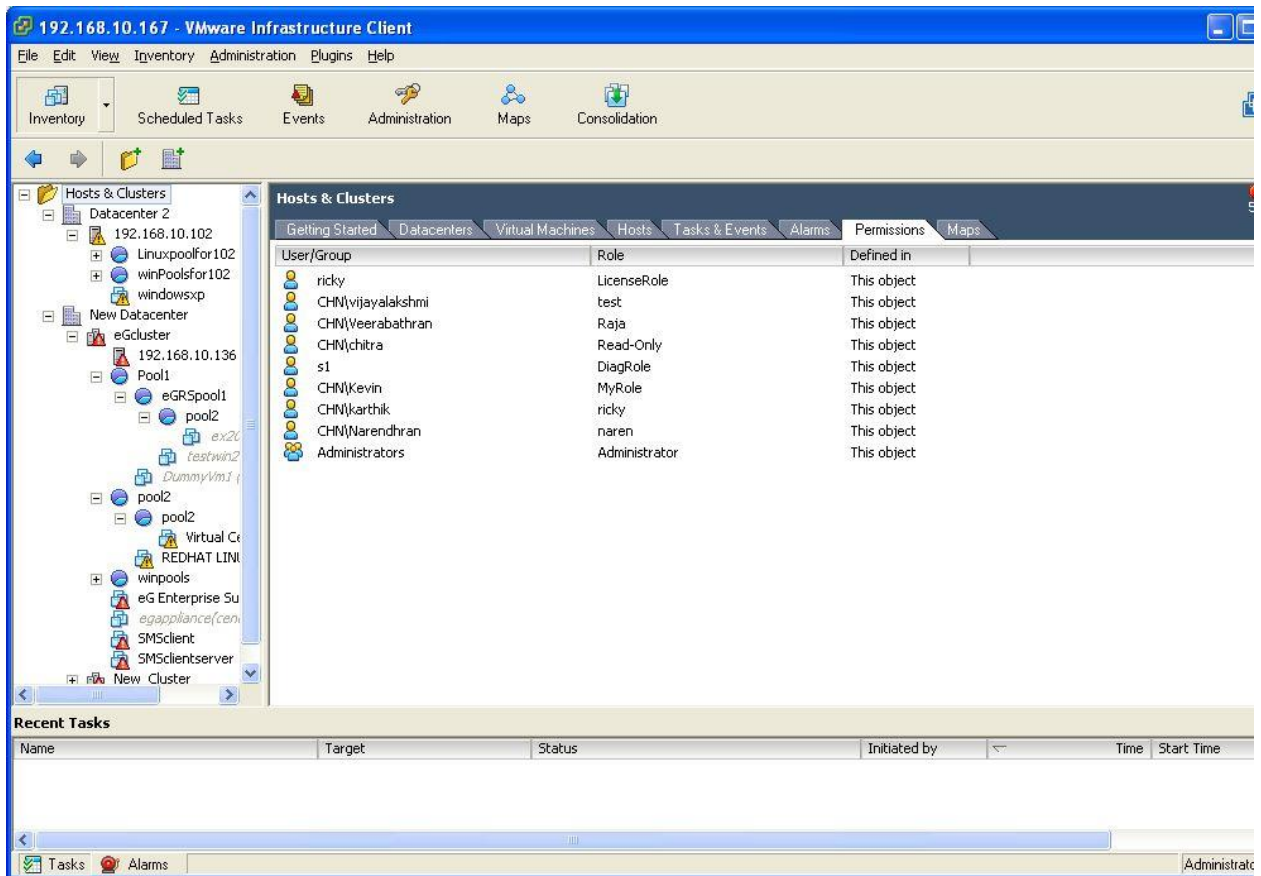


Figure 17: The Permissions tab page displaying the chosen user and the role assigned to him/her

## 1.1.6    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.

| Purpose | 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 vCenter server |
| Agent deploying the test | An internal agent |

| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed |
|---|---|
| | 2. **HOST** - The host for which the test is to be configured |
| | 3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 4. **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 Section 8.3 of 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 Section 8.4 of the *Monitoring VMware Infrastructures* document. |
| | 5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box. |
| | 6. **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. |
| | 7. **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. |

| | 8. **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 enabled/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. |
|---|---|
| **Outputs of the test** | One set of results for the vCenter server being monitored |

| Measurements made by the test | **Measurement** | **Measurement Unit** | **Interpretation** |
|---|---|---|---|
| | **ESX servers managed by VCenter:**<br><br>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:**<br><br>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:**<br><br>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:**<br><br>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:**<br><br>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. |

| | | | |
|---|---|---|---|
| | **ESX servers disconnected from VCenter:**<br><br>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:**<br><br>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:**<br><br>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.<br><br>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:**<br><br>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:**<br><br>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:<br><br>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.<br><br>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:**<br><br>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.<br><br>A virtual machine can become orphaned, in any of the following situations:<br><br>e.  After a VMotion or VMware DRS Migration;<br><br>f.  After a VMware HA host failure occurs or after the ESX Server host comes out of maintenance mode;<br><br>g.  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;<br><br>h.  If vCenter is restarted while a migration is in progress; this is a temporary situation though;<br><br>i.  If you schedule too many virtual machines to be relocated at the same time;<br><br>To know the names of the orphaned VMs, use the detailed diagnosis of this measure, if enabled. |

| | | | |
|---|---|---|---|
| | **Disconnected VMs managed by vCenter:**<br><br>Indicates the number of VMs that are currently disconnected from the ESX server. | Number | |
| | **Invalid VMs managed by vCenter:**<br><br>Indicates the number of VMs that are currently invalid. | Number | |
| | **Inaccessible VMs managed by vCenter:**<br><br>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 18: 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 19: The detailed diagnosis of the ESX servers connected to VC measure

## 1.1.7     vCenter User Sessions Test

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

| Purpose | Reports the number of sessions for every user who is currently logged into VC |
|---|---|
| Target of the test | A vCenter server |
| Agent deploying the test | An internal agent |

| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed |
|---|---|
| | 2. **HOST** - The host for which the test is to be configured |
| | 3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 4. **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 Section 8.3 of 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 Section 8.4 of the *Monitoring VMware Infrastructures* document. |
| | 5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box. |
| | 6. **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. |
| | 7. **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. |

| | 8. **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 enabled/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. |
| **Outputs of the test** | One set of results for the every currently logged in user to VC |
| **Measurements made by the test** | **Measurement** | **Measurement Unit** | **Interpretation** |
| | **Current sessions to VCenter:**<br><br>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.



Figure 20: The detailed diagnosis of the Current sessions to VC measure

## 1.1.8    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.
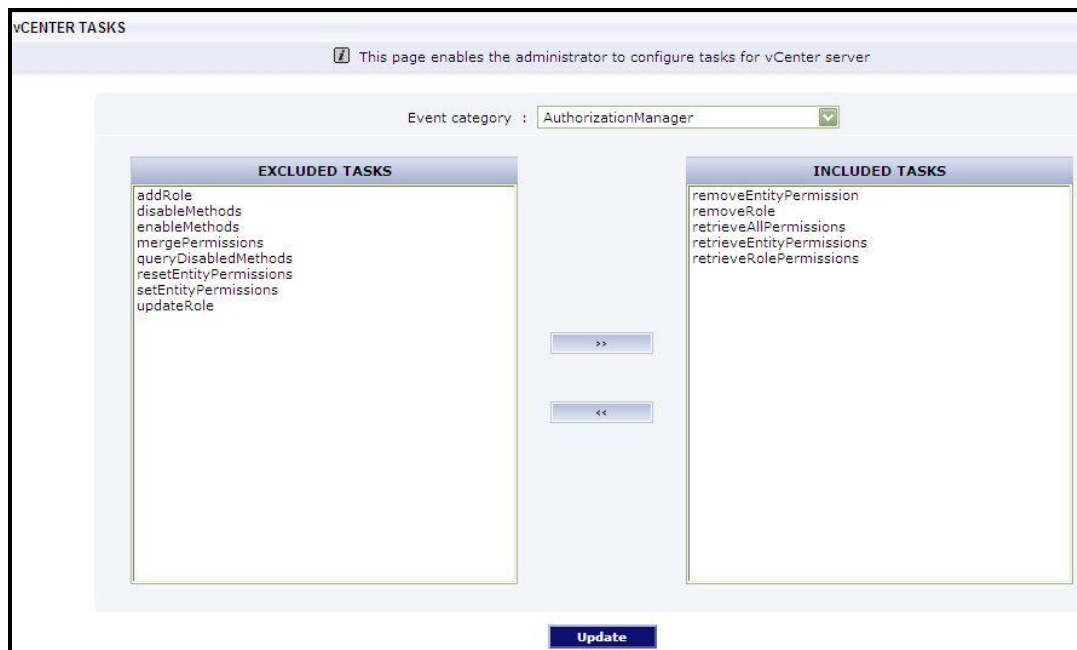
3.    Either way, Figure 26 appears.



Figure 21: Configuring the tasks to be monitored

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

**Note**:

By default, the **Event category** list in Figure 26 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.systemManagement*

➢ 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 of Figure 26**.

➢ 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:**

➤ 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.

---

**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**.

---

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.

| Purpose | Monitors configured VC tasks and reports their current status |
|---|---|
| **Target of the test** | A vCenter server |
| **Agent deploying the test** | An internal agent |

| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed |
|---|---|
| | 2. **HOST** - The host for which the test is to be configured |
| | 3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 4. **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 Section 8.3 of 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 Section 8.4 of the *Monitoring VMware Infrastructures* document. |
| | 5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box. |
| | 6. **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. |
| | 7. **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. |

8. **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.

9. **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**.

10. **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 enabled/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.

| Outputs of the test | One set of results for every event category that has been configured for monitoring | | |
|---|---|---|---|
| **Measurements made by the test** | **Measurement** | **Measurement Unit** | **Interpretation** |
| | **Started tasks:**<br><br>Indicates the number of tasks started during this measurement period.<br><br>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:**<br><br>Indicates the number of tasks that ended during this measurement period.<br><br>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:**<br><br>Indicates the number of tasks that succeeded during this measurement period.<br><br>For the *All* descriptor, this measure will report the total number of successful tasks across all event categories, during this measurement period. | Number | Use the detailed diagnosis of this measure to view the tasks that succeeded. |
| | **Failed tasks:**<br><br>Indicates the number of tasks that failed during this measurement period.<br><br>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.<br><br>Use the detailed diagnosis of this measure to view the tasks that failed. |
| | **Percent of failures:**<br><br>Indicates the percentage of tasks that failed.<br><br>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 | |
| | 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 22: The detailed diagnosis of the successful tasks measure

## 1.1.9    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.
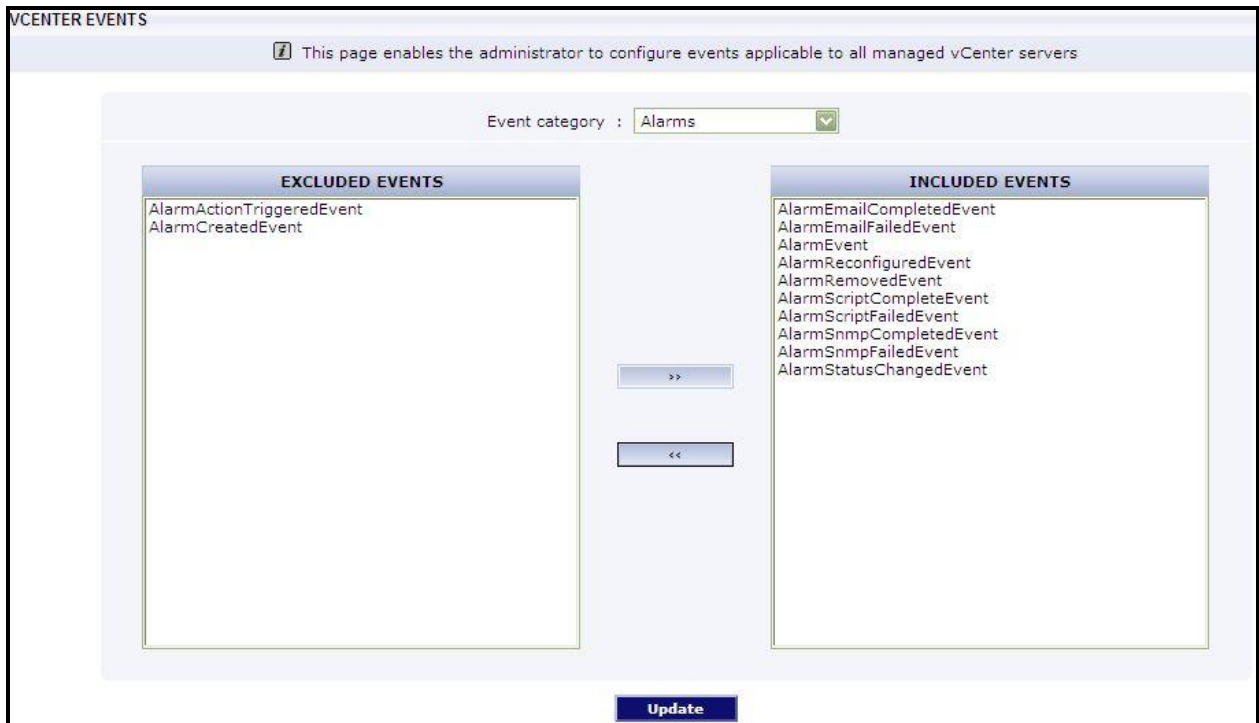
3.  Either way, Figure 26 appears.



Figure 23: Configuring the events to be monitored

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

**Note:**

By default, the **Event category** list in Figure 28 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 **[VirtualCenterEvents]** section of this file, you will find the default event categories that will populate the **Event category** list of Figure 28 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 of Figure 26. 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.randomAccessOpenReadonl y,NfcService.getVmFiles,NfcService.putVmFiles,NfcService.fileManagement,NfcService.systemMan agement*

➢ 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 of Figure 28.**

➢ 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:**

1.   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.

2.   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>*

3.   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*.

4.   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 of Figure 28, 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. To include any of the events that are by default listed in the **EXCLUDED EVENTS** list, select the event from the list and click the **<<** button. Similarly, to exclude any of the events that are by default available in the **INCLUDED EVENTS** list, select one/more events from the list and click the **>>** button.

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

| Purpose | Monitors configured events and reports their current status |
|---|---|
| Target of the test | A vCenter server |
| Agent deploying the test | An internal agent |

| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed |
|---|---|
| | 2. **HOST** - The host for which the test is to be configured |
| | 3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 4. **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 Section 8.3 of 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 Section 8.4 of the *Monitoring VMware Infrastructures* document. |
| | 5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box. |
| | 6. **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. |
| | 7. **SHOWALLEVENTS** – 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 events have registered for that measure. This enables administrators to easily assess the overall performance of events configured for monitoring on a vCenter server. |

8. **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.

9. **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 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.

10. **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**.

11. **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 enabled/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.

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

| Measurements made by the | Measurement | Measurement Unit | Interpretation |
|---|---|---|---|

| test | 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. |
|---|---|---|---|
| | 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:**<br><br>Indicates the number of user events of this category that occurred on vCenter in the last measurement period.<br><br>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.<br><br>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 24: 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 25: The detailed diagnosis of the Error measure

## 1.1.10    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.

The measures made by this test are as follows:

| Purpose | Ttracks the different types of virtual machine migrations occurring on vCenter, and reports their success/failure |
|---|---|
| **Target of the test** | A vCenter server |
| **Agent deploying the test** | An internal agent |

| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed |
|---|---|
| | 2. **HOST** - The host for which the test is to be configured |
| | 3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 4. **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 Section 8.3 of 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 Section 8.4 of the *Monitoring VMware Infrastructures* document. |
| | 5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box. |
| | 6. **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. |
| | 7. **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. |

| | |
|---|---|
| | 8. **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.<br><br>The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:<br><br>    ▪    The eG manager license should allow the detailed diagnosis capability<br><br>    ▪    Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. |
| **Outputs of the test** | One set of results the vCenter server being monitored |

| Measurements made by the test | Measurement | Measurement Unit | Interpretation |
|---|---|---|---|
| | **Total migrations:**<br><br>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:**<br><br>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:**<br><br>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. |
| | **DRS migrations:**<br><br>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:**<br><br>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:**<br><br>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:**<br><br>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. |

## 1.1.11    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. |

| **Purpose** | Enables administrators to better manage the vRAM and the vSphere 5 licenses by periodically reporting the vRAM allocation to VMs and their usage |
|---|---|

| Target of the test | A vCenter server that manages vSphere 5 servers |
|---|---|
| Agent deploying the test | An internal agent |
| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed<br><br>2. **HOST** - The host for which the test is to be configured<br><br>3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL.<br><br>4. **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 Section 8.3 of the *Monitoring VMware Infrastructures* document.<br><br>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 Section 8.4 of the *Monitoring VMware Infrastructures* document.<br><br>5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box.<br><br>6. **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.<br><br>7. **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.<br><br>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. |

| | | | |
|---|---|---|---|
| | 8. **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.<br><br>The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:<br><br>    ▪    The eG manager license should allow the detailed diagnosis capability<br><br>    ▪    Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| **Outputs of the test** | One set of results the vCenter server being monitored | | |
| **Measurements made by the test** | **Measurement** | **Measurement Unit** | **Interpretation** |
| | **Pooled vRAM capacity:**<br><br>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:**<br><br>Indicates the number of VMs that are currently powered on. | Number | |
| | **Current vRAM usage:**<br><br>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. |
| | **vRAM utilization:**<br><br>Indicates the percentage of total vRAM capacity that is currently utilized by the powered on VMs. | Percent | 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. |
| | **Total VMs:**<br><br>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 all the VMs are powered on. | GB | These metrics help administrators 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. |
| --- | --- | --- | --- |
| | 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 | 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:  j. Whether to procure additional vSphere licenses, so that the aggregate vRAM capacity increases, or,  k. Whether to power on only a few VMs at a time, so that the vRAM availability is not badly impacted; |
| | 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. |

## 1.1.12    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 component-type for which these tests are to be enabled as the **Component type**, set *Performance* as the **Test type**, choose the tests from the **DISABLED TESTS** list, and click on the **>>** button to move the tests to the **ENABLED TESTS** list. Finally, click the **Update** button.

| Purpose | Promptly alerts administrators to the addition, removal, and migration of VMs managed by vCenter |
| --- | --- |
| Target of the test | A vCenter server that manages vSphere 5 servers |
| Agent deploying the test | An internal agent |

| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed |
|---|---|
| | 2. **HOST** - The host for which the test is to be configured |
| | 3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 4. **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 Section 8.3 of 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 Section 8.4 of the *Monitoring VMware Infrastructures* document. |
| | 5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box. |
| | 6. **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. |
| | 7. **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. |

| | | | |
|---|---|---|---|
| | 8. **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 enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <br> ▪ The eG manager license should allow the detailed diagnosis capability <br> ▪ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. | | |
| **Outputs of the test** | One set of results for the vCenter server being monitored | | |
| **Measurements made by the test** | **Measurement** | **Measurement Unit** | **Interpretation** |
| | **VMs created:** <br><br> 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: <br><br> • The name of the VM that was added; <br><br> • The host to which the VM was added; <br><br> • The cluster to which the host belongs; <br><br> • The resource pool to which the VM belongs; <br><br> • The CPU, memory, and disk space configuration of the VM; <br><br> • The datastore assigned to the VM; |

| | | Removed VMs: | 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: |
|---|---|---|---|---|
| | | Indicates the number of VMs that were removed from vCenter during the last measurement period. | | • The name of the VM that was removed; <br><br> • The host from which the VM was removed; <br><br> • The cluster to which the host belongs; <br><br> • The resource pool from which the VM was removed; <br><br> • The CPU, memory, and disk space configuration of the removed VM; <br><br> • The datastore that was assigned to the VM; |
| | | **Total migrations:** <br><br> 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: <br><br> • The name of the VM that was migrated; <br><br> • The host from which the VM was migrated; <br><br> • The host to which the VM was migrated; <br><br> • The cluster from which the VM was migrated; <br><br> • The cluster to which the VM was migrated; <br><br> • The resource pool from which the VM was migrated; <br><br> • The resource pool to which the VM was migrated; <br><br> • The CPU, memory, and disk space configuration of the migrated VM; <br><br> • The datastore assigned to the VM; |

# 1.1.13   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.
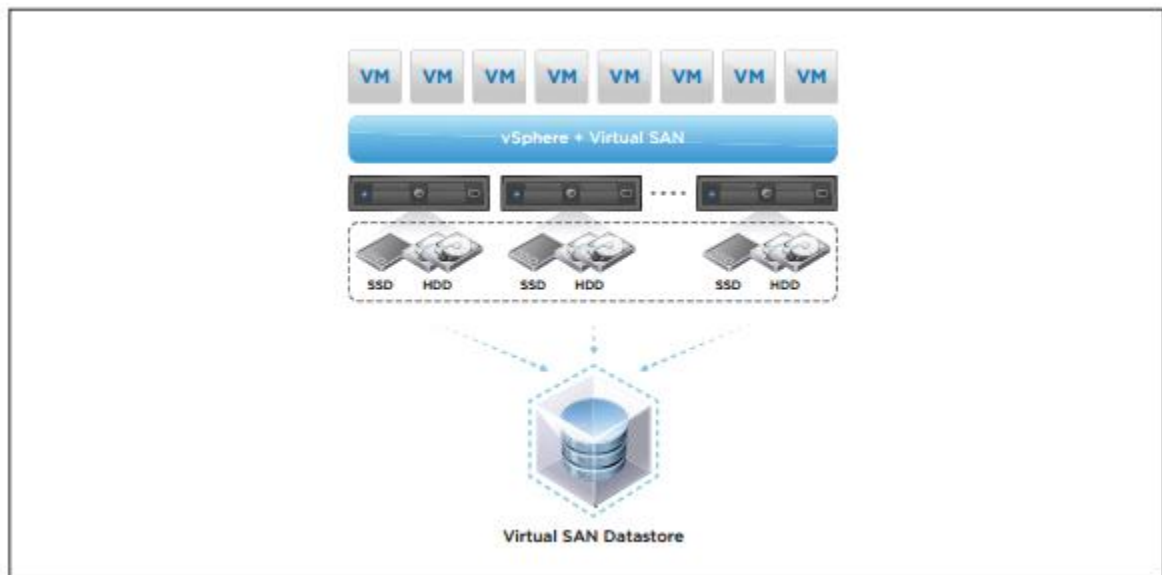


Figure 3.1: VMware virtual SAN

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 is 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.

| Purpose | 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 | An ESX server host |
| Agent deploying the test | An internal/remote agent |

| Configurable parameters for the test | 9. **TEST PERIOD** - How often should the test be executed |
|---|---|
| | 10. **HOST** - The host for which the test is to be configured |
| | 11. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 12. **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 Section 8.3 of 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 Section 8.4 of the *Monitoring VMware Infrastructures* document. |
| | 13. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box. |
| | 14. **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. |
| | 15. **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. |
| Outputs of the test | One set of results for every vSphere/ESX server host that is participating in the vSAN service |
| Measurements made by the | **Measurement** | **Measurement Unit** | **Interpretation** |

| test | **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:<br><br>| **Measure Value** | **Numeric Value** |<br>| --- | --- |<br>| Healthy | 1 |<br>| Unhealthy | 0 |<br><br>**Note**:<br><br>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. |
| --- | --- | --- | --- |

## 1.2 The Virtual Machine Clusters Layer

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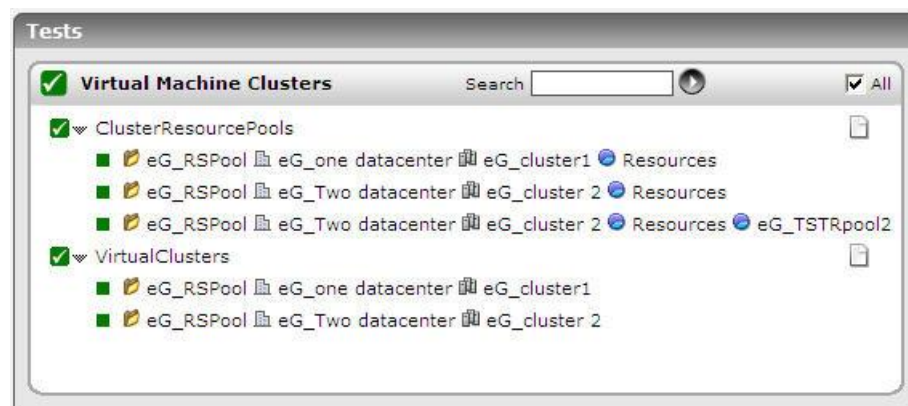


Figure 26: The tests mapped to the vCenter Cluster layer

### 1.2.1    Virtual Clusters Test

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

| Purpose | Reports key metrics pertaining to the resource availability and resource usage of the ESX server clusters managed by vCenter |
| --- | --- |
| Target of the test | A vCenter server |
| Agent deploying the test | An internal agent |

| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed |
|---|---|
| | 2. **HOST** - The host for which the test is to be configured |
| | 3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 4. **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 Section 8.3 of 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 Section 8.4 of the *Monitoring VMware Infrastructures* document. |
| | 5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box. |
| | 6. **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. |
| | 7. **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. |

a.

| 8. | **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.<br><br>The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:<br><br>▪ The eG manager license should allow the detailed diagnosis capability<br><br>▪ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. |
| --- | --- |

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

| **Measurements made by the test** | **Measurement** | **Measurement Unit** | **Interpretation** |
| --- | --- | --- | --- |
| | **Total CPU in cluster:**<br><br>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. |
| | **CPU available to cluster:**<br><br>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 amount of resources available for the root resource pool for running virtual machines. |
| | **Physical CPU usage of VMs in cluster:**<br><br>Indicates the current physical CPU usage of the VMs in the cluster in Mhz. | Mhz | |

| | Physical CPU used by VMs:<br><br>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.<br><br>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. |
| --- | --- | --- | --- |
| | Total memory in cluster:<br><br>Indicates the aggregated memory resources of all hosts in the cluster. | MB | |
| | Memory available to cluster:<br><br>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. |
| | Host memory used by VMs in cluster:<br><br>Indicates the amount of host memory currently used by the VMs in the cluster. | MB | |
| | Physical memory used by cluster:<br><br>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:**<br><br>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:**<br><br>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 cluster. |
| | **Number of effective hosts in cluster:**<br><br>Indicates the total number of effective hosts in the cluster, currently. | Number | |
| | **VMs in cluster:**<br><br>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:**<br><br>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:**<br><br>Indicates the total number of VMs that are currently powered-off in the cluster. | Number | |
| | **VMs suspended:**<br><br>Indicates the number of VMs that are currently in a suspended state in the cluster. | Number | |
| | **Effective CPU reserved:**<br><br>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:**<br><br>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:**<br><br>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 cluster:**<br><br>Indicates the total amount of CPU resources reserved by the direct pools in the cluster. | Mhz | Use the detailed diagnosis of this measure to figure out how much CPU has been reserved by each of the direct pools in the cluster. |
| | **CPU used by direct pools in cluster:**<br><br>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:**<br><br>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:**<br><br>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:**<br><br>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 capability of this measure to quickly zero-in on that VM that consumes maximum memory. |
| | **Total migrations by VMotion:** Indicates the total number of migrations with VMotion that have been done internal to this cluster. | Number | |
| | **Usage of effective CPU by cluster:** Indicates the percentage of effective CPU resources currently available to run the virtual machines. | Percent | A high value is desired for this measure. |
| | **Usage of effective memory:** Indicates the percentage of effective memory resources currently available to run the virtual machines. | Percent | A high value is desired for this measure. |

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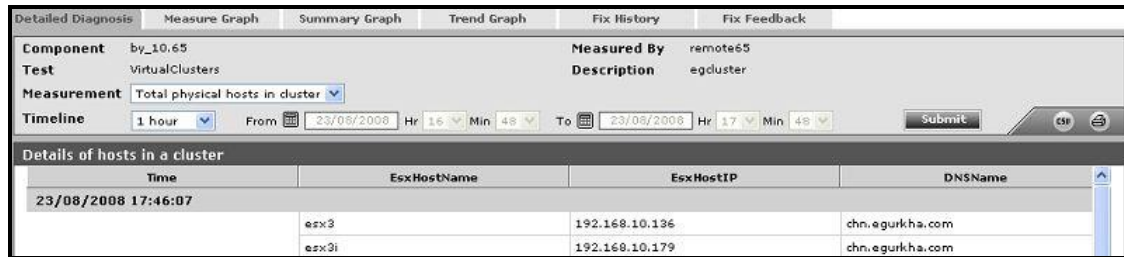
Figure 27: The detailed diagnosis of the Total physical hosts in cluster measure

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 IP address of the ESX server.
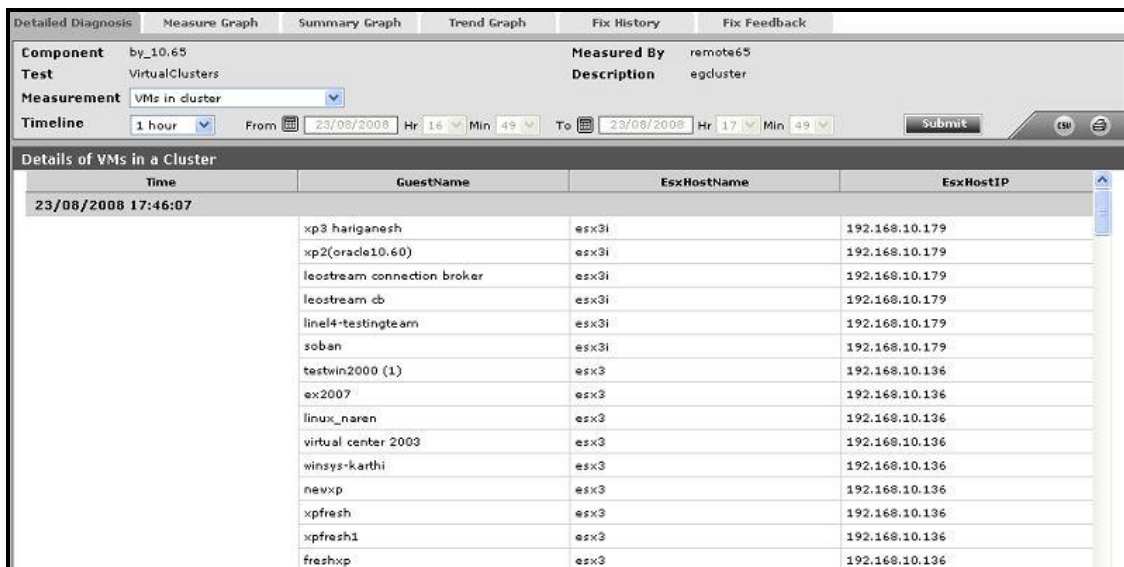


Figure 28: The detailed diagnosis of the VMs in cluster measure

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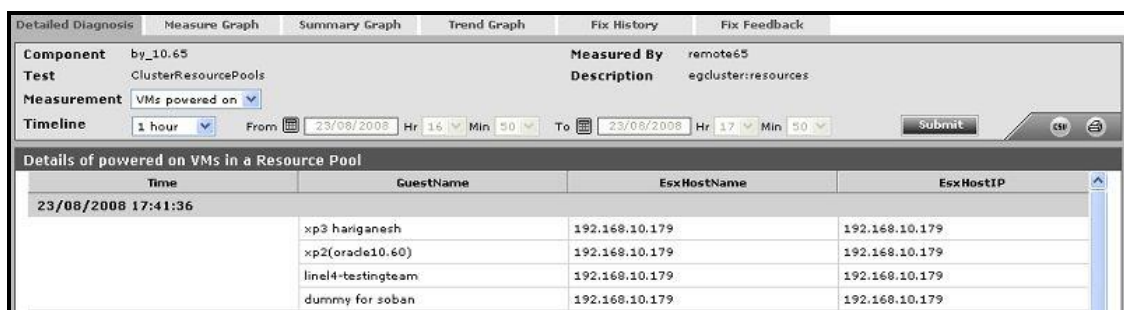


Figure 29: The detailed diagnosis of the VMs powered on measure

## 1.2.2   Cluster Resource Pools Test

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

| Purpose | Auto-discovers the resource pools and/or vApps under ESX clusters, and reports critical statistics pertaining to each pool |
|---------|---------------------------------------------------------------------------------------------------------------------------|

| Target of the test | A 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 | 1. **TEST PERIOD** - How often should the test be executed<br><br>2. **HOST** - The host for which the test is to be configured<br><br>3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL.<br><br>4. **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 Section 8.3 of the *Monitoring VMware Infrastructures* document.<br><br>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 Section 8.4 of the *Monitoring VMware Infrastructures* document.<br><br>5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box.<br><br>6. **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.<br><br>7. **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.<br><br>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. |

8. **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 enabled/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 | Measurement Unit | Interpretation |
|---|---|---|---|
| | **Pool status:** <br><br> Indicates the current status of this resource pool. | | 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. <br><br> The numeric values that correspond to these measures are as follows: <br><br> <table><tr><td>**State**</td><td>**Value**</td></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> |
| | **Child resource pools:** <br><br> Indicates the total number of child resource pools under this resource pool. | Number | |
| | **VMs in pool:** <br><br> Indicates the total number of VMs in this resource pool. | Number | This number includes both the directly and indirectly assigned VMs. |

| | | | |
|---|---|---|---|
| | **VMs powered on:**<br><br>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:**<br><br>Indicates the number of virtual machines that were directly assigned to this resource pool. | Number | |
| | **Direct VMs powered on:**<br><br>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:**<br><br>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:**<br><br>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.<br><br>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 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:**<br><br>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.<br><br>Follow the guidelines mentioned below to define the **Reservation** setting:<br><br>➢ 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.<br>➢ 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.<br>➢ If you expect frequent changes to the total available resources, use **Shares**, not **Reservation**, to allocate resources fairly across resource pools.<br><br>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:**<br><br>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.<br><br>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 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 VMs in this resource pool. | MB | |

| | **Memory capacity:**<br><br>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.<br><br>Follow the guidelines mentioned below to define the **Reservation** setting:<br><br>➢ 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.<br>➢ 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.<br>➢ If you expect frequent changes to the total available resources, use **Shares**, not **Reservation**, to allocate resources fairly across resource pools.<br><br>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 capacity reserved:**<br><br>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.<br><br>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:**<br><br>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.<br><br>If the resource pool does not use the reserved memory, then remember the following:<br><br>   ➢ Other resource pools *cannot reserve* these memory resources<br><br>   ➢ Other resource pools *can use* these memory resources, that is, idle memory reservations are not wasted. |
| | **Free memory in pool:**<br><br>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:**<br><br>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.<br><br>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:**<br><br>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:**<br><br>Indicates the percentage of total configured or available memory in this pool. | Percent | |

| | **Zero memory:**<br>Indicates the amount of memory that is zeroed out. | MB | |
|---|---|---|---|
| | **Swap memory:**<br>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:<br><br>  ➢  Was never installed<br>  ➢  Has been explicitly disabled<br>  ➢  Is not running (for example, while the guest operating system is booting)<br>  ➢  Is temporarily unable to reclaim memory quickly enough to satisfy current system demands<br><br>Standard demand-paging techniques swap pages back in when the virtual machine needs them.<br><br>Swap space must be reserved on disk 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.<br><br>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:<br><br>&#10148; The additional time to access memory within a virtual machine.<br>&#10148; The extra space needed by the ESX Server host for its own code and data structures, beyond the memory allocated to each virtual machine.<br><br>The ESX server's memory virtualization ensures that the time overhead to memory accesses is minimal.<br><br>The memory space overhead, on the other hand, is composed of two other components:<br><br>&#10148; A fixed system-wide overhead for the service console (in the case of ESX 3/3.5) and the VMkernel.<br>&#10148; Additional overhead for each virtual machine<br><br>In addition, the space reserved for the virtual machine frame buffer and various virtualization data structures, also add to the memory overhead.<br><br>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:**<br><br>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.<br><br>The maximum amount of memory that can be reclaimed from a guest may be configured by modifying the "sched.mem.maxmemctl" advanced option.<br><br>If the memory reclaimed from a guest (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:**<br><br>Indicates the amount of memory in the resource pool that is shared. | MB | |
| | **Memory granted:**<br><br>Indicates the amount of memory that is granted to this resource pool. | MB | |

# 1.3 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 30: The tests mapped to the Datacenters layer

## 1.3.1    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.

| Purpose | 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 | vCenter |
| Agent deploying the test | An internal agent |

| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed |
|---|---|
| | 2. **HOST** - The host for which the test is to be configured |
| | 3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 4. **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 Section 8.3 of 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 Section 8.4 of 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**. |
| | 5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box. |
| | 6. **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. |

7. **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.

8. **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 enabled/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.

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

| **Measurements made by the test** | **Measurement** | **Measurement Unit** | **Interpretation** |
|---|---|---|---|
| | **Esx servers using the datastore:**<br><br>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.<br><br>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:**<br><br>Indicates the number of VMs that are currently using this datastore. | Number | |
| | **Physical disk capacity:**<br><br>Indicates the total capacity of the datastore. | MB | |
| | **Used space:**<br><br>Indicates the space in the datastore that is currently in use. | MB | |

| | | | |
|---|---|---|---|
| | **Free space:**<br><br>Indicates the space currently available for use in the datastore. | MB | |
| | **Physical disk space usage:**<br><br>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:**<br><br>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.<br><br>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:**<br><br>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:**<br><br>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 requests rate:**<br><br>Indicates the average number of read commands issued per second to this datastore. | Commands/Sec | |
| | **Write latency:**<br><br>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:**<br><br>Indicates the average amount of time taken to read from this datastore. | Secs | |

| | | | |
|---|---|---|---|
| | **Read rate:**<br><br>Indicates the rate at which data is read from this datastore. | MB/Sec | |
| | **Write rate:**<br><br>Indicates the rate at which data is written to this datastore. | MB/Sec | |
| | **Storage I/O control normalized latency:**<br><br>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.<br><br>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:**<br><br>Indicates the aggregate number of I/O operations that are occurring on this datastore. | Number | |
| | **Storage I/O Control active time:**<br><br>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.<br><br>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:**<br><br>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:**<br><br>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. Enabling IOC 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. |
| | **Throughput:**<br><br>Indicates the rate at which the data is read and written to this datastore. | MB/Sec | 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. |

| | | | The values that this measure can report are listed below, along with their descriptions: |
|---|---|---|---|

**Type:**

Indicates the type of file system this datastore supports.

| Measure Value | Description |
|---|---|
| VMFS | 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. |
| 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 | ? |
| VFAT | Virtual File Allocation Table (VFAT) handles long file 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 into pools of storage that can be carved up and assigned to virtual |

95

| Measure Value | Description |
|---|---|
| VFFS | 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. |

The numeric values that correspond to each of the **Measure Value**s listed above are as follows:

| Measure Value | Numeric Value |
|---|---|
| VMFS | 1 |
| NFS | 2 |
| NFSV41 | 3 |
| CIFS | 4 |
| VFAT | 5 |
| VSAN | 6 |
| VFFS | 7 |

**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 31: 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 32: The detailed diagnosis of the VMs using the datastore measure

## 1.3.2    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.

| Purpose | 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 | vCenter |
| Agent deploying the test | An internal agent |

| Configurable parameters for the test | 1. **TEST PERIOD** - How often should the test be executed |
|---|---|
| | 2. **HOST** - The host for which the test is to be configured |
| | 3. **PORT –** Refers to the port at which vCenter listens. By default, this is NULL. |
| | 4. **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 Section 8.3 of 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 Section 8.4 of the *Monitoring VMware Infrastructures* document. |
| | 5. **CONFIRM PASSWORD** – Confirm the password by retyping it in this text box. |
| | 6. **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. |

b.

<table>
<tr>
<td rowspan="2"></td>
<td colspan="3">7. <strong>WEBPORT</strong> - 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 <strong>SSL</strong> flag above is set to <strong>No</strong>), then the eG agent connects to vCenter using port 80 by default, and if vCenter is SSL-enabled ((i.e., if the <strong>SSL</strong> flag is set to <strong>Yes</strong>), then the agent-vCenter communication occurs via port 443 by default.  Accordingly, the <strong>WEBPORT</strong> parameter is set to <em>default</em> by default.

In some environments however, the default ports 80 or 443 might not apply. In such a case, against the <strong>WEBPORT</strong> 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.

8. <strong>DETAILED DIAGNOSIS</strong> - 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 <strong>On</strong> option. To disable the capability, click on the <strong>Off</strong> option.

The option to selectively enabled/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

<ul>
<li>The eG manager license should allow the detailed diagnosis capability</li>
<li>Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li>
</ul></td>
</tr>
<tr></tr>
<tr>
<td><strong>Outputs of the test</strong></td>
<td colspan="3">One set of results for the every <em>datacenter:datastore</em> combination being monitored</td>
</tr>
<tr>
<td rowspan="4"><strong>Measurements made by the test</strong></td>
<td><strong>Measurement</strong></td>
<td><strong>Measurement Unit</strong></td>
<td><strong>Interpretation</strong></td>
</tr>
<tr>
<td><strong>ESX servers in datacenter:</strong>

Indicates the number of ESX servers in this datacenter.</td>
<td>Number</td>
<td></td>
</tr>
<tr>
<td><strong>ESX servers directly assigned to datacenter:</strong>

Indicates the number of ESX servers that are directly connected to the datacenter.</td>
<td>Number</td>
<td>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.</td>
</tr>
<tr>
<td><strong>ESX servers under clusters in datacenter:</strong>

Indicates the number of ESX servers in the datacenter that are under clusters.</td>
<td>Number</td>
<td>The detailed diagnosis of this measure, if enabled, lists the ESX servers that are under clusters.</td>
</tr>
</table>

| | VMs in datacenter: | Number | |
|---|---|---|---|
| | Indicates the number of VMs in the datacenter. | | |
| | **VMs directly assigned to servers in datacenter:**<br><br>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:**<br><br>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:**<br><br>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:**<br><br>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:**<br><br>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 33: 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 34: 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 35: The detailed diagnosis of the VMs under cluster measure

# 1.4 Troubleshooting vCenter Appliance Monitoring

The VMware vCenter Server Appliance is a preconfigured Linux-based virtual machine that is optimized for running vCenter Server and associated services. When monitoring this appliance in an agent-based manner, the eG agent will fail to start if the *eG user* does not have shell access. Likewise, when monitoring this appliance in an agentless manner, the eG remote agent will not be able to report metrics for the **Operating System** layer, if the SSH user configured for agentless monitoring does not have shell access.

To resolve this issue, you should grant shell access to the *eG user* or the SSH user, depending upon how the vCenter appliance is monitored (i.e., whether in an agent-based or an agentless manner).

To grant shell access to the *eG user* in case of agent-based monitoring, do the following:

- Login to the vCenter appliance as a *root* user.

- Open the shell prompt.

- Run the **vipw** command. This command is used to edit the file **/etc/passwd**.

- If the command runs successfully, the contents of the **/etc/passwd** file will be displayed in editable mode.

- In the contents, look for an entry that begins with the name of the *eG user*. For instance, if the

user who started the eG agent on the vCenter appliance is *john*, then look for an entry that begins with *john*. Typically, this entry will be as follows:

*john:x:1012:4045::/opt/egurkha:/bin/appliancesh*

Change this entry as follows:

*john:x:1012:4045::/opt/egurkha:/bin/bash*

- Finally, save the file.

To grant shell access to the SSH user who has been configured for monitoring the vCenter appliance in an agentless manner, do the following:

- First, login to the eG management console as *admin*, and find out what value you specified for the **User** parameter when managing the VMware vCenter appliance in an agentless manner. For this, select the *Add/Modify* option from the *Components* menu of the *Infrastructure* tile. Pick *VMware vCenter* as the **Component type** to view all the vCenter servers that are being managed by eG. In this list, locate the vCenter appliance that has been managed in an agentless manner, and click the **Modify** icon corresponding to it. When the details of the appliance are displayed, look for the value of the **User** parameter, and make a note of it.

- Then, login to the vCenter appliance as a *root* user.

- Open the shell prompt.

- Run the **vipw** command. This command is used to edit the file **/etc/passwd**.

- If the command runs successfully, the contents of the **/etc/passwd** file will be displayed in editable mode.

- In the contents, look for an entry that begins with the user name you noted in step 1. For instance, if the remote agent connects to and monitors the vCenter appliance as user *elvis*, then look for an entry that begins with *elvis*. Typically, this entry will be as follows:

*elvis:x:1012:4045::/opt/egurkha:/bin/appliancesh*

Change this entry as follows:

*elvis:x:1012:4045::/opt/egurkha:/bin/bash*

- Finally, save the file.

# 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.
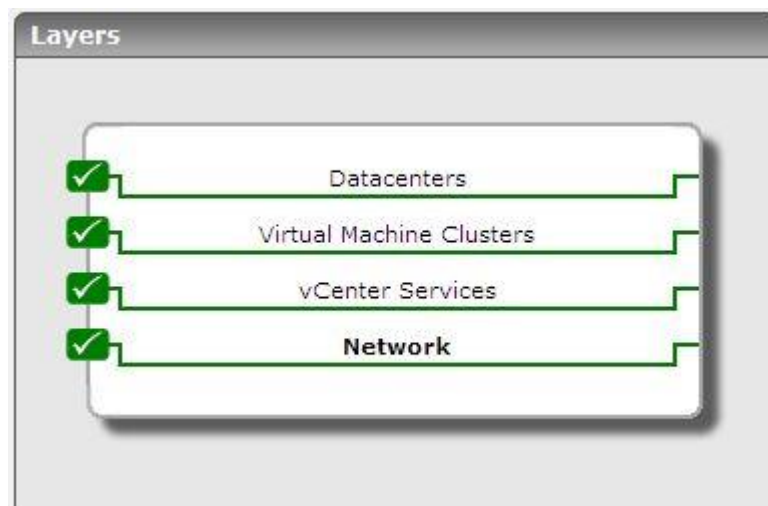


Figure 36: The VMware vCenter Cluster monitoring model

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.

Chapter

# 3

# Conclusion

This document has described in detail the monitoring paradigm used and the measurement capabilities of the eG Enterprise suite of products with respect to the **VMware vCenter**. For details of how to administer and use the eG Enterprise suite of products, refer to the user manuals.

We will be adding new measurement capabilities into the future versions of the eG Enterprise suite. If you can identify new capabilities that you would like us to incorporate in the eG Enterprise suite of products, please contact support@eginnovations.com. We look forward to your support and cooperation. Any feedback regarding this manual or any other aspects of the eG Enterprise suite can be forwarded to feedback@eginnovations.com.