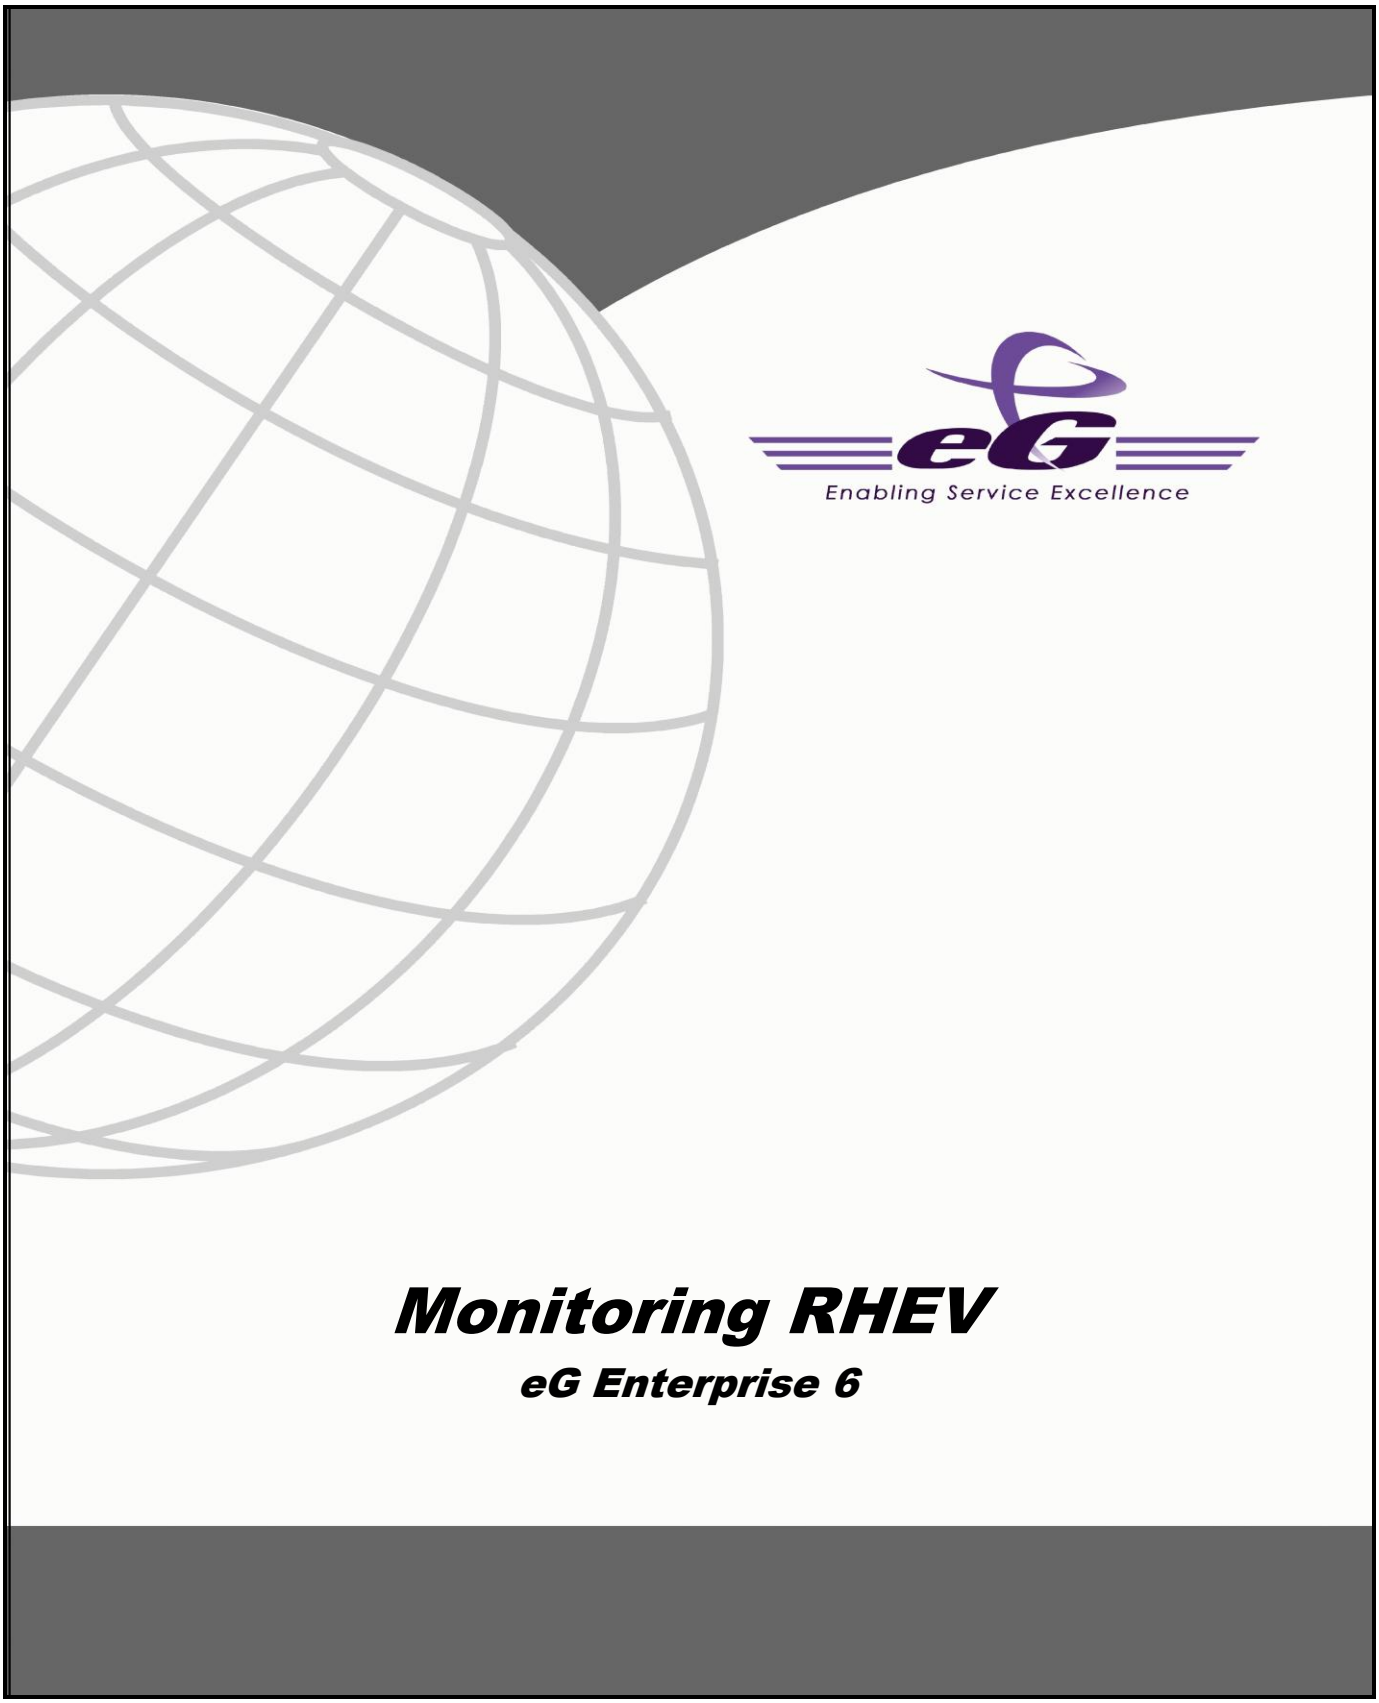

Monitoring RHEV

eG Enterprise 6

Restricted Rights Legend

Trademarks

Copyright

Table of Contents

Table of Figures

Chapter 1

Introduction

Red Hat Enterprise Virtualization is an end-to-end virtualization solution with use cases for both servers and desktops. The Red Hat Enterprise Virtualization portfolio provides IT departments with the tools to meet the challenges of managing complex environments. This state-of-the-art virtualization suite enables administrators to reduce the cost and complexity of large deployments, for example, for thousands of virtual machines.

The suite consists of the following components:

- **Red Hat Enterprise Virtualization Hypervisor**, which is a thin virtualization layer deployed across the server's infrastructure.
- **Agents and tools** include VDSM, which runs in the hypervisor or host. These provide local management for virtual machines, networks and storage.
- **Red Hat Enterprise Virtualization platform management infrastructure** allows users to view and manage all the system components, machines and images from a single, powerful interface.

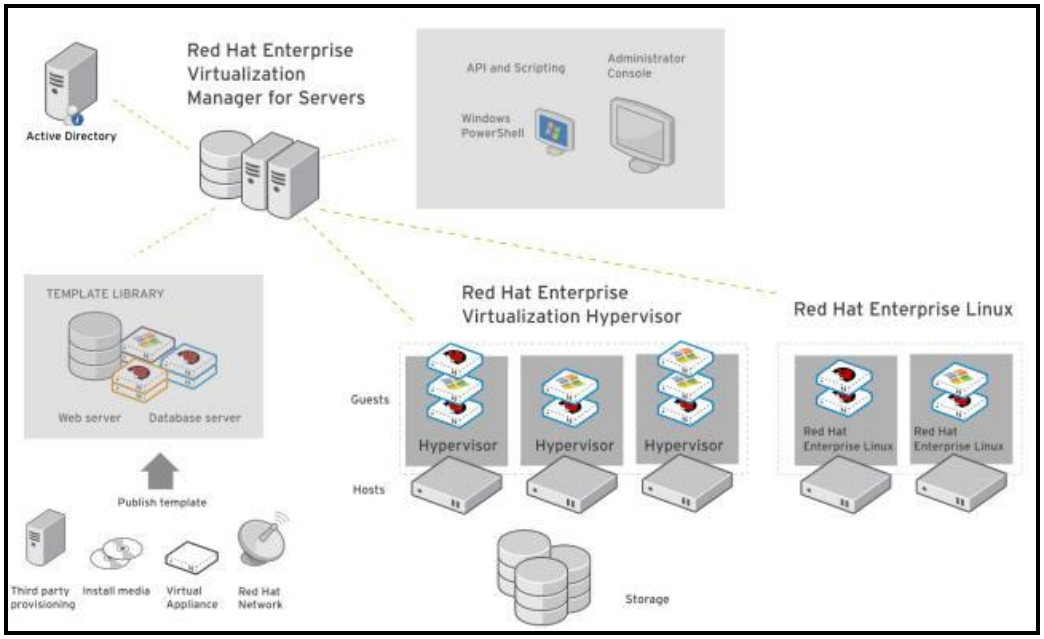

Figure 1.1: The core components of the RHEV portfolio

The **Red Hat Enterprise Virtualization (RHEV) Hypervisor** is the core component of the RHEV suite. It is a compact, full-featured virtualization platform for quickly and easily deploying and managing virtualized guests. It is designed for integration with the Red Hat Enterprise Virtualization Manager for Servers and the Red Hat Enterprise Virtualization Manager for Desktops - this implies that you can deploy mission-critical server applications (e.g., Web server, Oracle database server, etc.) or simple desktop applications (e.g., Microsoft Word, Adobe Acrobat Reader/Writer, Notepad, etc.) on the VMs on the RHEV hypervisor.

This hypervisor is based on the Kernel-based Virtual Machine (KVM). KVM is an advanced and efficient virtualization hypervisor implemented as a Linux kernel module. As KVM is a kernel module, it leverages the existing Red Hat Enterprise Linux kernel and benefits from the default kernel's extensive testing, device support and flexibility.

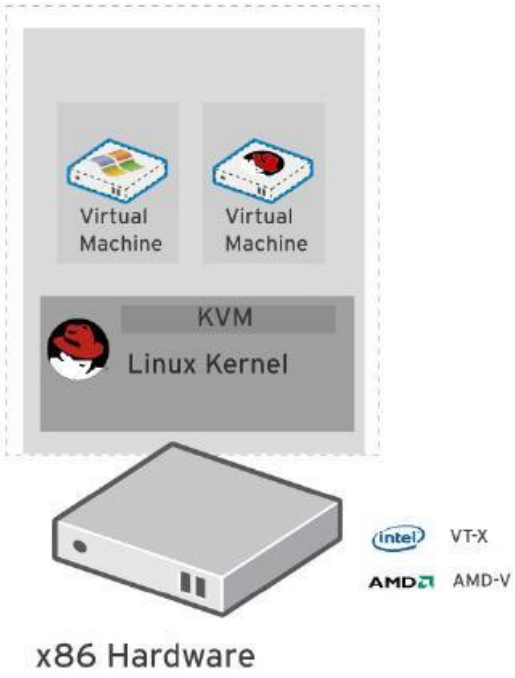

Figure 1.2: Architecture of the RHEV Hypervisor

1.1 Challenges in Monitoring RHEV

While RHEV eases virtualization adoption and improves operational efficiency at one end, on the other, it introduces new dynamic dependencies - between physical machines and VMs, between applications and VMs, and between the VMs on a hypervisor; these dependencies are difficult to comprehend, manage, and monitor! A single malfunctioning application on a VM can therefore degrade the performance of the applications running on other VMs on the same RHEV hypervisor. To compound the problem, VMs on an RHEV Hypervisor can be migrated from one physical machine to another - under such circumstances, it is almost herculean to keep track of the movement of the VMs and assess their performance impact on each physical host they visit.

The desktop virtualization deployment model of RHEV also poses myriad challenges! For instance, while the server application virtualization approach typically involves a smaller number of VMs running

on a physical server, in the virtual desktop approach, tens of desktop VMs run on a physical server. The scale of deployment of desktops makes monitoring and management a pain!

While in-depth monitoring of each of the applications is important in the server application virtualization approach, in the virtual desktop approach, since only client applications are executed on the desktop, monitoring of the desktop need not be as in-depth as in the server application virtualization context. Furthermore, in a virtual desktop environment, it is essential to identify which guest a user is logging on to, for how long the user was logged in, and what applications he/she used. This information is critical for planning the capacity of the virtual desktop environment. Notice should also be taken of the fact that in a virtual desktop environment, virtual desktops may come and go off dynamically (e.g., as a user logs on and logs off, respectively), whereas in a server application virtualization approach, the guest operating systems are likely to be more static (i.e., come on and off less frequently).

To ensure business continuity and user satisfaction in such virtualized environments, you need a solution that can:

- Auto-discover the VMs on an RHEV hypervisor and identify the applications operating on each VM;
- Assess the resource usage of the hypervisor host and the VMs, rapidly isolate resource-starved VMs, and precisely point to the application/process on the VM that is causing the resource drain;
- Automatically determine physical machine - VM, VM - application, VM - VM relationships, intelligently correlate performance in the light of these dependencies, and accurately diagnose root-cause of performance issues;
- Understand the unique monitoring needs of virtual desktop environments, and automatically shift the monitoring focus from 'VMs' to 'Users';
- Instantly detect and report the migration of a VM and continuously track the performance ramifications of the migration on each of the destination hosts;

1.2 How does eG Enterprise Monitor RHEV?

The eG Monitor for RHEV, part of the eG Enterprise Suite, is a comprehensive solution for monitoring and managing all aspects of virtual hosts and guests, whether the infrastructure is used to support server or desktop applications. Toward this end, the eG Enterprise Suite offers two specialized monitoring models - the *RHEV Hypervisor* model that focuses on the server virtualization platform and the *RHEV Hypervisor - VDI* model that monitors the desktop virtualization platform.

Both these models extend the universal monitor technology of the eG Enterprise Suite to virtualized environments. Using a patented **In-N-Out Monitoring** approach, these models provide a comprehensive view of the RHEV server, including the performance of the hypervisor and all of its VMs.

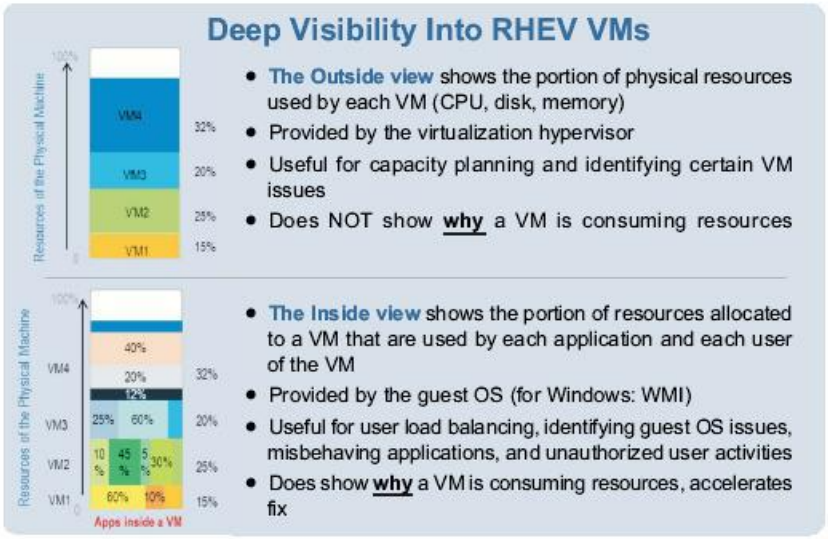

Figure 1.3: The In-N-Out view

An eG agent deployed on a remote Windows or a Linux system connects to the RHEV manager that is managing the RHEV Hypervisor to be monitored via HTTP/HTTPS, auto-discovers the IP address and operating system of the VMs running on that hypervisor, and uses the RHEV RESTful APIs to provide an *outside view* of every VM's performance. The relative resource usage levels of the VMs show where the performance hogs may lie (see Figure 1.4).

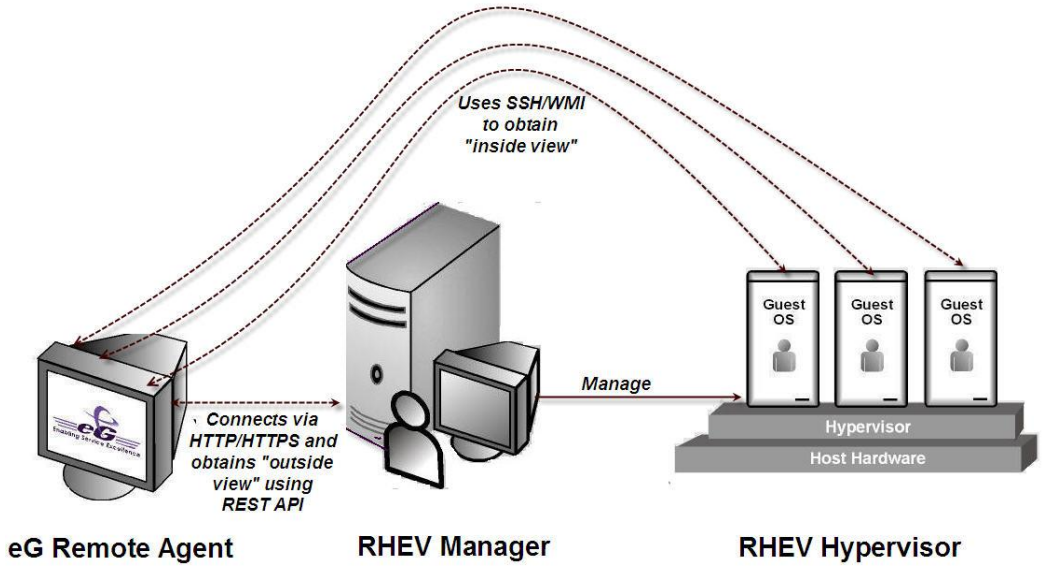

Figure 1.4: Agentless monitoring of an RHEV Hypervisor

To complement the outside view, the eG agent obtains an "inside view" that details the user activity, resource allocation and the application mix inside the VMs. To obtain this "inside view", the eG agent connects to every VM on the monitored RHEV Hypervisor via SSH/WMI and pulls out the detailed

metrics. To establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM, which allows the eG agent on the service console to collect "inside view" metrics from the VMs **without domain administrator rights.** Refer to Section 1.4 for more details on the **eG VM Agent**.

To obtain the inside and outside views, the pre-requisites detailed in Section 1.3 are to be fulfilled.

1.3 Pre-requisites for Monitoring RHEV

1.3.1 General Pre-requisites

- Enable the remote agent to communicate with the eG manager port (default: 7077).
- If VMs running on multi-byte operating systems are to be monitored (eg., *Windows Japanese*), then the remote agent monitoring such VMs should also run on a multi-byte operating system.
- 32-bit VMs that are to be monitored in an agentless manner should be configured with at least 2 GB RAM, and 64-bit VMs require at least 4 GB RAM. If more than four RHEV Hypervisors are being monitored in an agentless manner, then the RAM capacity of the VMs should be increased proportionately.

1.3.2 Pre-requisites for Obtaining the "Outside View" of VMs by connecting to the RHEV Manager

- Ensure that the remote agent has IP connectivity to the RHEV manager.
- Ensure that the remote agent has HTTP/HTTPS access to the port (port 8080/8443) at which the RHEV manager listens
- All the tests that the remote agent executes should be configured with the name and password of a user with *read-only access* to the REST API of the RHEV manager. To know how to configure a *read-only* role on the RHEV manager and assign that role to a user, follow the steps detailed in Section 1.4 of this document.

1.3.3 Pre-requisites for Obtaining the "Inside View" of Windows VMs, using the eG VM Agent

- Install the eG VM Agent on each Windows VM. For details on how to install the eG VM Agent, refer to Section 1.4 of this document.
- Enable the remote agent to communicate with the port at which the eG VM Agent listens (default port: 60001).
- Set the **INSIDE VIEW USING** flag for all the "inside view" tests to **eG VM Agent (Windows)**.

1.3.4 Pre-requisites for Obtaining the "Inside View" of VMs, without using the eG VM Agent

- Ensure that the remote agent has IP connectivity to at least one of the network interfaces of the VMs.
- The **ADMIN$** share should be enabled for all Windows-based virtual guests being monitored and the administrative account must have permissions to this share drive. Refer to Section 1.6.1 of this document for a step-by-step procedure to achieve this.
- To enable the remote agent to communicate with the Windows VMs, an administrative account login and password (either a local account or a domain account) must be provided when configuring the eG monitoring capabilities.
- In case of VMs with the Windows XP/Windows 2003/Windows 2008/Windows Vista/Windows 7 operating systems, the firewall on the guest should be explicitly configured to allow Windows File and Print Sharing services which are required for the remote agent to communicate with the guest operating system.
- For monitoring a Windows VM, TCP port 139 must be accessible from the remote agent to the VM.
- For monitoring a Linux VM, the SSH port (TCP port 22) must be enabled for communication between the remote agent and the VM being monitored.

Note:

If the Linux VMs in your environment listen on a different SSH port, then, you can override the default SSH port of 22 using the steps provided below:

- Login to the eG manager.
- Edit the **eg_tests.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory) on the eG manager host.
- In the **[AGENT_SETTINGS]** section of the file, set the **JavaSshPortForVm** parameter to an SSH port of your choice. By default, this parameter is set to *22*.
- If your environment consists of multiple Linux VMs, each listening on a different SSH port, then, you can specify a comma-separated list of SSH ports against the **JavaSshPortForVm** parameter. For example: *7711,7271,8102*
- Finally, save the file.

- For obtaining the "inside view" of VMs running Windows Vista/Windows 7/Windows 2008 operating systems, the **eGurkhaAgent** service of the eG remote agent should be configured to run using *domain administrator* privileges. Refer to the *eG User Manual* for the procedure. For obtaining the "inside view" of other Windows VMs however, the remote agent service requires no such privileges.
- Set the **INSIDE VIEW USING** flag for all the "inside view" tests to **Remote connection to VM (Windows)**.

1.4 Configuring the eG Agent to use the RESTful APIs on the RHEV Manager to Obtain the "Outside View"

The eG agent uses the RESTful APIs on the RHEV manager to report the *outside view* of performance of the VMs on the RHEV hypervisor. To be able to connect to the RESTful API, the eG agent should be configured with the credentials of a user with *read-only* access to the API. To create a *read-only* role and assign it to a user, do the following:

1. Open the RHEV manager.
2. In the top navigation bar, click **Configure**. The configuration window opens.
3. In the configuration window, click the **Roles** tab.
   - To create a role, click **New**.
   - In the **New Role** window, provide the name of the role. Select **Admin** as the account type and leave all the check boxes in the **Check Boxes to Allow Action** pane clear. Click **OK**.
4. In the configuration window, click the **System Permission** tab.
   - To grant a user with the permission to access information about the virtual machines, click **Add**.
   - In the **Add System Permission to User** pane, select the user to whom you want to grant the permission.
   - From the **Assign role to user** list, select the role that you created in step 3 and click **OK**.

1.5 Configuring the Remote Agent to Obtain the Inside View of Windows VMs, using the eG VM Agent

To provide the inside view of a Linux VM, the eG agent uses secure shell (SSH). To obtain the inside view of a Windows VM, the eG agent offers two options. The first option uses Windows File & Print Sharing services to push monitoring components to the VMs. These monitoring components are then executed on the VM to collect metrics from the VMs. To push monitoring components to the VM and to periodically invoke these components, the eG agent requires **domain administrator privileges** to all the VMs being monitored.

In many production environments, strict security restrictions are enforced, and it may not be possible to configure a monitoring solution with domain administration privileges for each of the VMs. To handle such environments, the eG RHEV monitor uses a lightweight monitoring component called the **eG VM Agent**, which is installed inside each of the VMs to obtain metrics regarding the health of the VMs. The **eG VM Agent** can be best described as a software that can be installed on the Windows virtual machines of a virtual infrastructure to allow a single eG agent to obtain an inside view of these VMs, **without domain administrator privileges**.

Users have multiple options to choose from when it comes to installing the eG VM Agent. These options have been discussed below:

- Manually install the eG VM Agent on every Windows VM using the executable that eG Enterprise includes;
- Bundle the eG VM Agent as part of a template VM, and use this template to create multiple VMs; this way, the eG VM Agent is automatically available in all the VMs that are created using the template;

- Use a software distribution solution such as Microsoft System Center to distribute the eG VM Agent software to existing VMs from a central location;

Use the install procedure that is ideal for your environment, and quickly get the eG VM Agent up and running. The detailed manual installation procedure has been discussed hereunder:

1. To install the eG VM Agent on a 32-bit VM, double-click on the **eGVMAgent.exe**, and to install the same on a 64-bit VM, double-click the **eGVMAgent_64.exe**.
2. Figure 1.5 then appears. Click on the **Next** button in Figure 1.5 to continue.

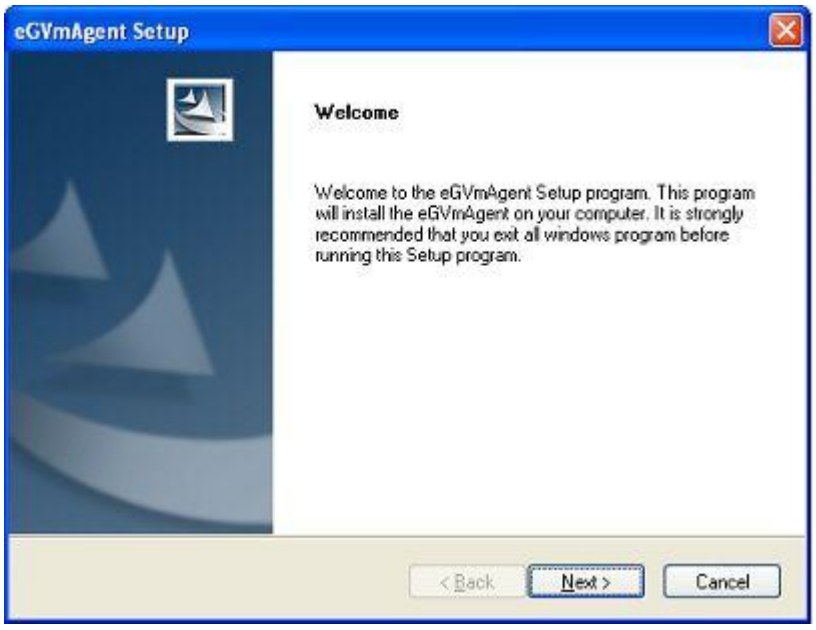

Figure 1.5: Welcome screen of the eG VM Agent installation wizard

3. When Figure 1.6 appears, click on **Yes** to accept the displayed license agreement.

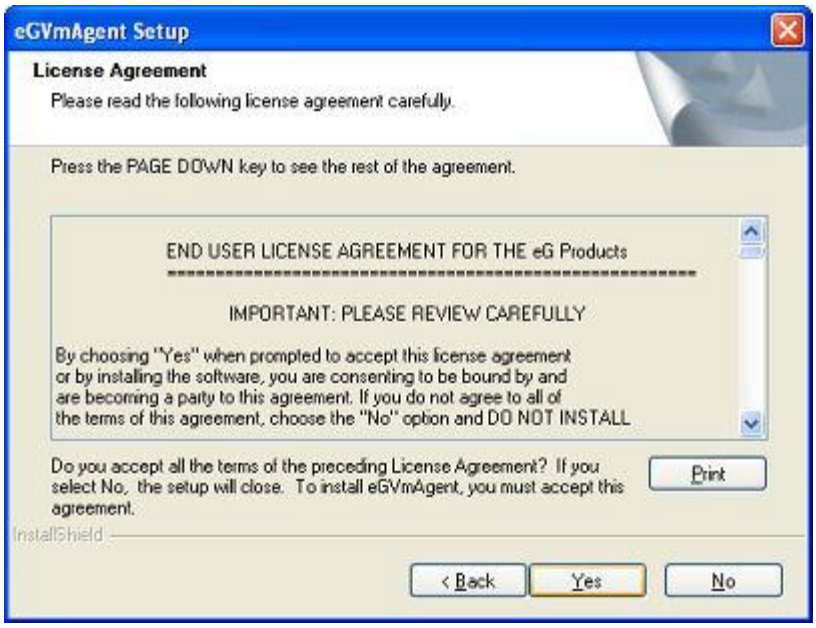

Figure 1.6: Accepting the license agreement

4. Use the **Browse** button in Figure 1.7 to indicate the location in which the agent should be installed, and click the **Next** button to proceed.

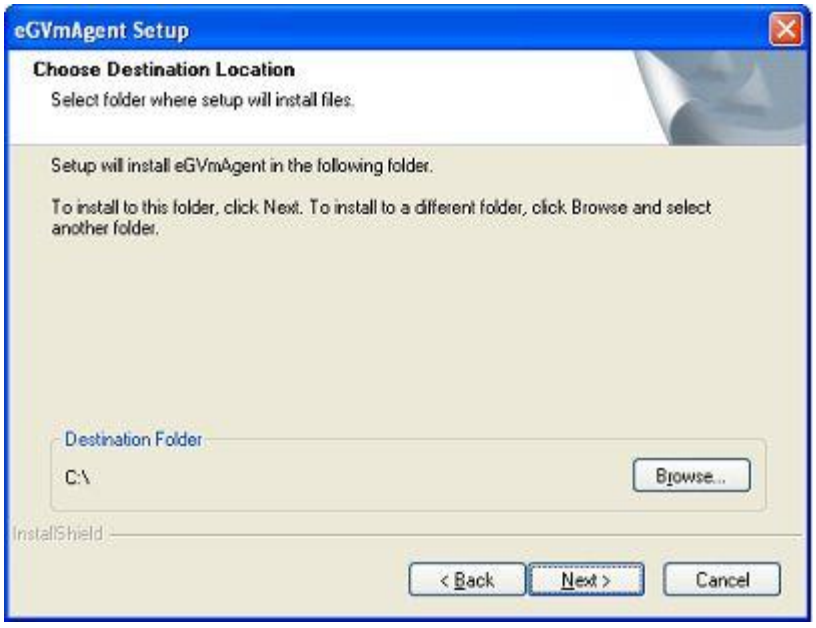

Figure 1.7: Specifying the install directory of the eG VM Agent

5. Next, specify the port at which the VM agent listens for requests from the eG agent. The default port is 60001. After port specification, click on the **Next** button in Figure 1.8 to proceed.

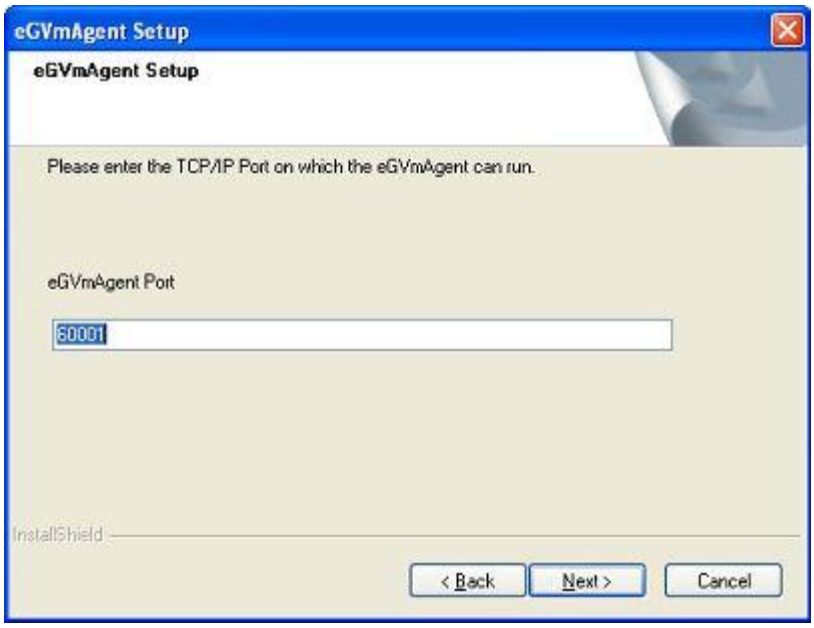

Figure 1.8: Specifying the VM agent port

6. A summary of your specifications then follows (see Figure 1.9). Click **Next** to proceed.

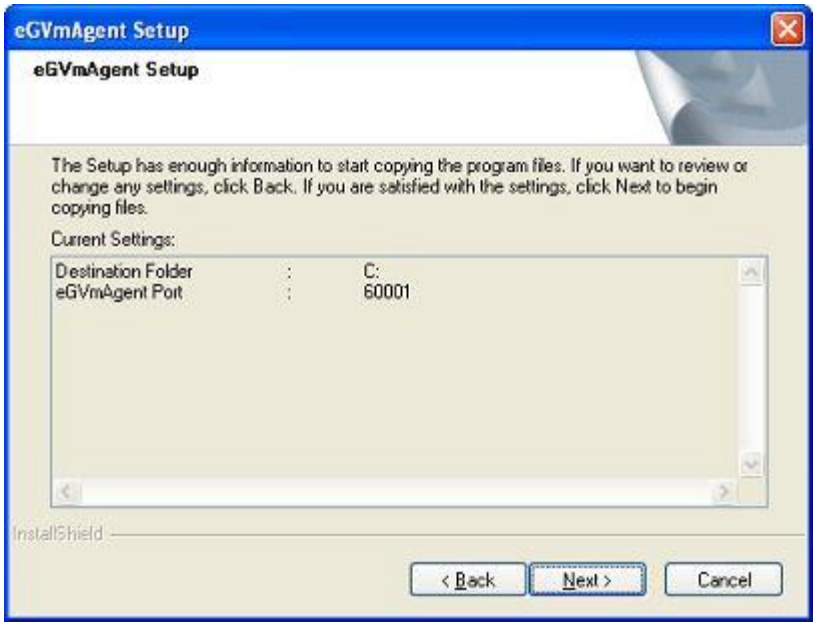

Figure 1.9: A summary of your specifications

7. Finally, click the **Finish** button in Figure 1.10 to complete the installation.

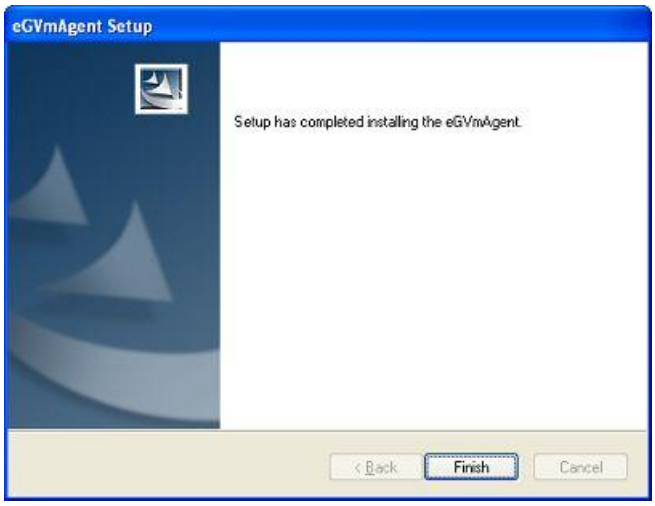

Figure 1.10: Finishing the installation

1.5.1 Communication between the eG Agent and the eG VM Agent

At the time of the installation of the eG VM agent, a folder named **eGVMAgent** is created in the install destination specified. The setup program also creates a Windows Service named **eGVMAgent** on the Windows VM. This service must be running for the eG agent to obtain the inside view of the virtual machine.

Upon successful installation, the eG VM agent starts automatically and begins listening for requests at default TCP port 60001. However, if, during the installation process, you have configured a different port for the eG VM agent, then, after completing the installation, follow the steps below to make sure that the eG agent communicates with the eG VM agent via the port that you have configured:

- Login to the eG manager host.
- Edit the **eg_tests.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- The **WmiInsideViewPort** parameter in the **[AGENT_SETTINGS]** section of the file is set to **60001** by default. If the eG VM agent's port is changed at the time of installation, then you will have to ensure that this parameter reflects the new port. Therefore, change the default port specification accordingly.
- Save the file.

At configured intervals, the eG remote agent issues commands to each of the eG VM Agents (using the TCP port configured during the VM agent installation). The eG VM Agent executes the commands, collects the "inside view" metrics from the Windows VM, and sends the output back to the eG agent. The eG agent then analyzes the metrics and informs the eG manager about the status of the Windows VMs.

1.5.2 Licensing of the eG VM Agent

The eG VM Agent is not license-controlled. Therefore, you can install and use any number of VM agents in your infrastructure.

1.5.3 Benefits of the eG VM Agent

The eG VM Agent offers several key benefits:

- **Ideal for high-security environments:** The eG VM Agent is capable of collecting "inside view" metrics from Windows VMs, without domain administrator privileges. It is hence ideal for high-security environments, where administrators might not be willing to expose the credentials of the domain administrators.
- **Easy to install, configure**: The eG RHEV Monitor offers users the flexibility to choose from multiple methodologies for installing the eG VM Agent on the target VMs. Even a manual installation procedure, would not take more than a few minutes. Moreover, since the eG VM agent communicates only with the eG agent and not the eG manager, no additional configuration needs to be performed on the VM agent to facilitate the communication. In addition, the VM agent starts automatically upon installation, thereby saving the time and trouble involved in manually starting each of the VM agents.
- **License independent:** Since the eG VM agent is not license-controlled, you can add any number of VM agents, as and when required, to your environment.

1.6 Configuring Windows Virtual Machines to Support the eG Agent's Inside View without the eG VM Agent

For the "inside" view, by default, the eG agent uses SSH/WMI (depending upon the virtual OS to be monitored) to communicate remotely with the virtual machines on the RHEV server and collect metrics. To establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. Besides, the **INSIDE VIEW USING** flag of all "inside view" tests should be set to **Remote connection to a VM**.

In addition, the following pre-requisites need to be fulfilled:

- The **ADMIN$** share will have to be available on the Windows guests
- The Windows Firewall should be configured to allow Windows File and Print Sharing

The sections to come discuss the procedure to be followed for fulfilling the 2 requirements above.

1.6.1 Enabling ADMIN$ Share Access on Windows Virtual Guests

1.6.1.1 Enabling ADMIN$ Share Access on Windows 2000/2003 VMs

If the **ADMIN$** share is not available on any Windows-based virtual guest, create the share using the procedure detailed below:

1. Open the Windows Explorer on the virtual machine, browse for the corresponding **Windows** directory in the C drive, right-click on it, and select the **Sharing** option from the shortcut menu.
2. If the **ADMIN$** share does not pre-exist on the Windows guest, then Figure 1.11 appears indicating the same.

Figure 1.11: The ADMIN$ share does not exist

On the other hand, if the **ADMIN$** share pre-exists, Figure 1.12 appears. In such a case, first, remove the **ADMIN$** share by selecting the **Do not share this folder** option from Figure 1.12 and clicking the **Apply** and **OK** buttons. After this, you will have to repeat step 1 of this procedure to open Figure 1.11. Then, proceed as indicated by step 3 onwards.

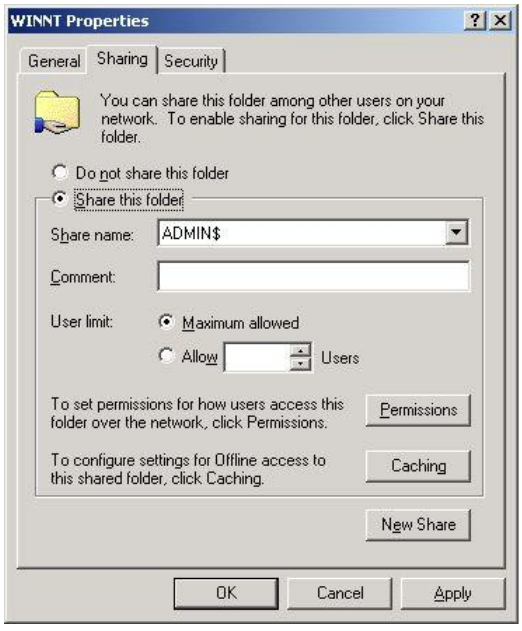

Figure 1.12: Admin$ share pre-exists

3. To create (or re-create) the **ADMIN$** share, select the **Share this folder** option from Figure 1.13, and provide **ADMIN$** share against the **Share name** text box (see Figure 1.13).

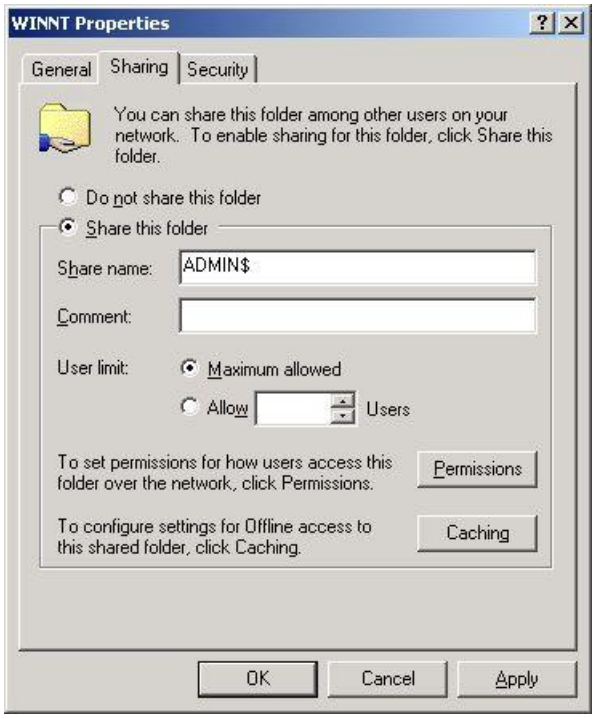

Figure 1.13: Creating the ADMIN$ share

4. Next, to enable the eG agent to communicate effectively with the Windows guest, you need to ensure that the permission to access the **ADMIN$** share is granted to an administrative user (local/domain); also, the **credentials of this user should be passed while configuring the eG monitoring capabilities** - i.e., while configuring the VMware tests. To grant the access permissions, click on the **Permissions** button in Figure 1.13.

5. By default, the **ADMIN$** share can be accessed by **Everyone** (see Figure 1.14). To grant access rights to a specific administrative (local/domain) user, select the **Add** button in Figure 1.14. When Figure 1.15 appears, select the domain to search from the **Look in** list. The valid user accounts configured on the chosen domain then appear in the box below. From this box, choose the administrator's account and click on the **Add** button to add the chosen user account to the box below the **Add** button.

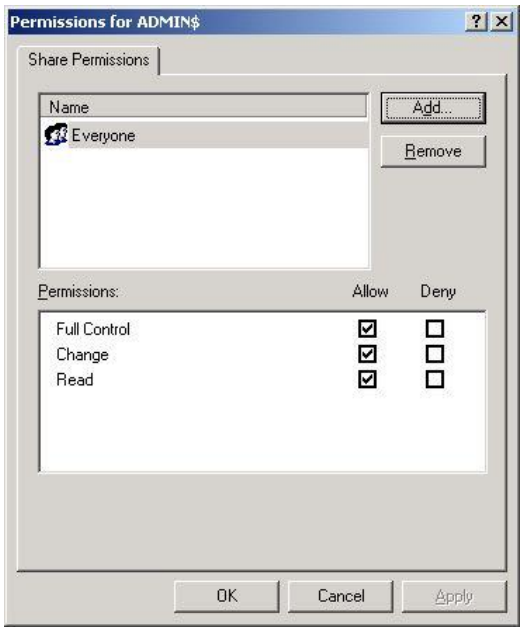

Figure 1.14: Clicking the Add button

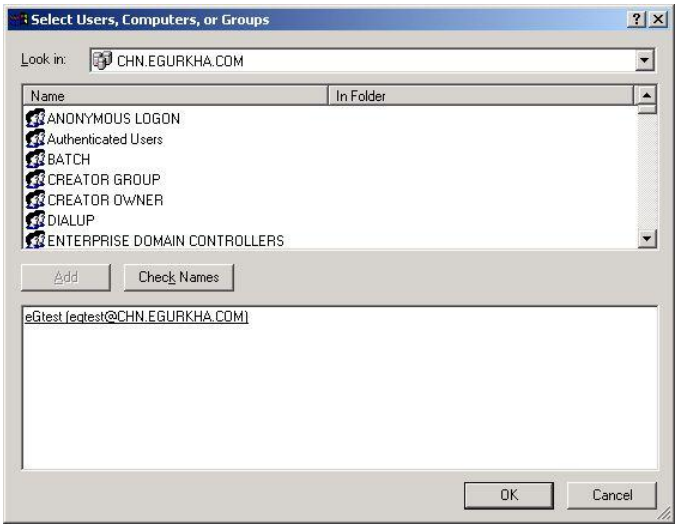

Figure 1.15: Selecting the administrative user to whom access rights are to be granted

6. Finally, click the **OK** button. You will then switch to Figure 1.16, where the newly added administrator account will appear.

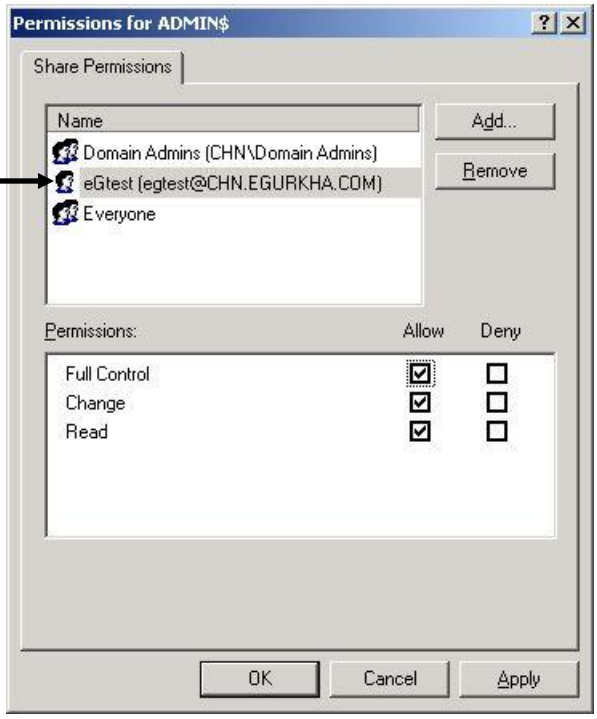

Figure 1.16: The administrator account granted access permissions

7. Select the newly added administrator account from Figure 1.16, and then, using the **Permissions** section, grant the administrator **Full Control**, **Change**, and **Read** permissions.
8. Finally, click the **Apply** and **OK** buttons in Figure 1.16 to register the changes.
9. Once you return to the **Properties** window, click on the **Security** tab to define the security settings for the **ADMIN$** share (see Figure 1.17).

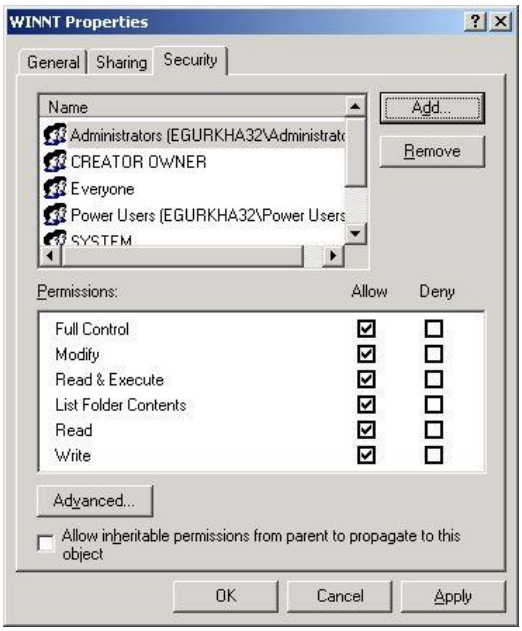

Figure 1.17: Defining the Security settings for the ADMIN$ share

10. Here again, you need to add the same administrator account, which was granted access permissions earlier. To do so, click the **Add** button in Figure 1.17, pick a domain from the **Look in** list of Figure 1.18, select the said administrator account from the domain users list below, and click the **Add** button (in Figure 1.18) to add the chosen account. Then, click the **OK** button in Figure 1.18.

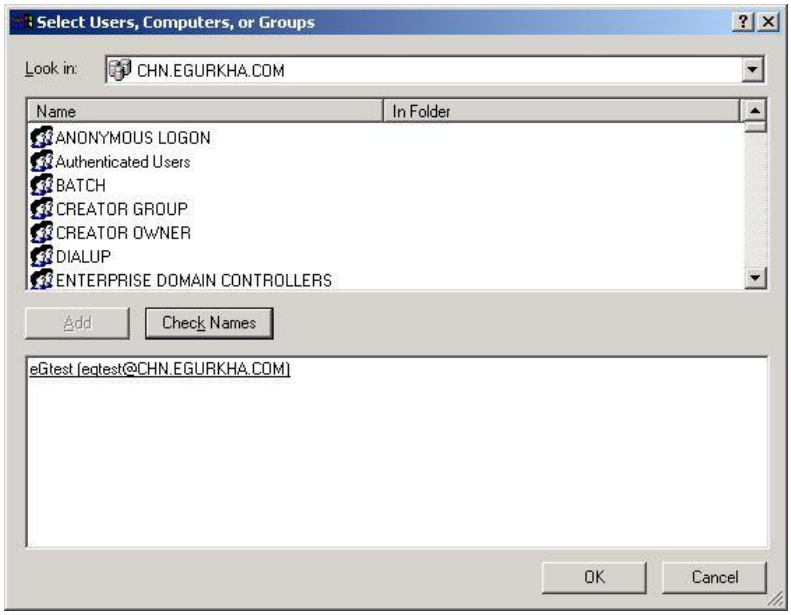

Figure 1.18: Adding the administrator account

11. This will bring you back to Figure 1.17, but this time, the newly added domain administrator account will be listed therein as indicated by Figure 1.19.

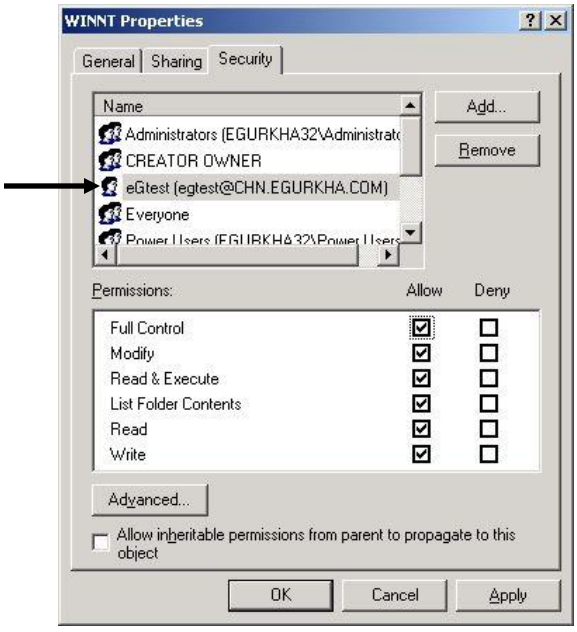

Figure 1.19: The Administrator account in the Security list

12. Finally, click the **Apply** and **OK** buttons in Figure 1.19.

1.6.1.2 Enabling ADMIN$ Share Access on Windows 2008 VMs

To enable the **ADMIN$** share on a Windows 2008 VM, do the following:

1. Open the Windows Explorer on the virtual machine, browse for the corresponding **Windows** directory in the C drive, right-click on it, and select the **Share** option from the shortcut menu.

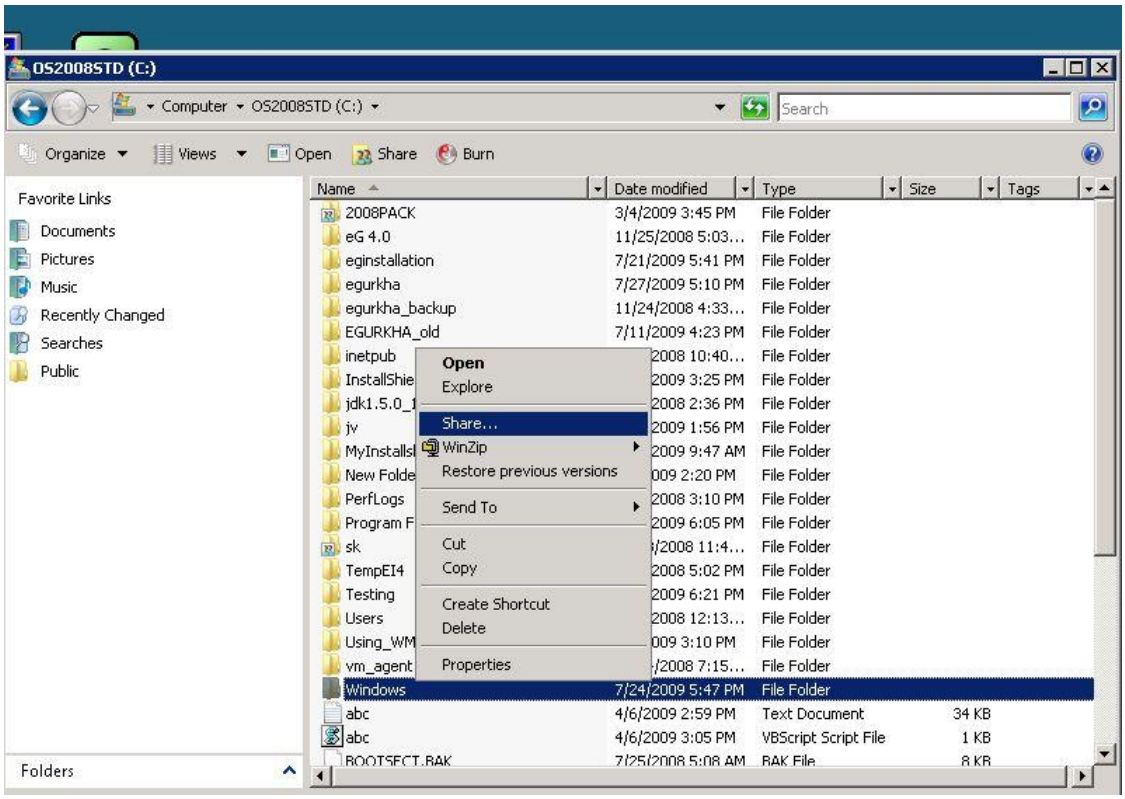

Figure 1.20: Selecting the Share option from the shortcut menu

8. Figure 1.21 will then appear. Click on **Advanced Sharing** in Figure 1.21.

Figure 1.21: Cicking on Advanced Sharing

9. Select the **Share this folder** check box in Figure 1.22 that appears, enter **ADMIN$** against **Share name**, and click on the **Permissions** button in Figure 1.22, to allow only a local/domain administrator to access the folder.

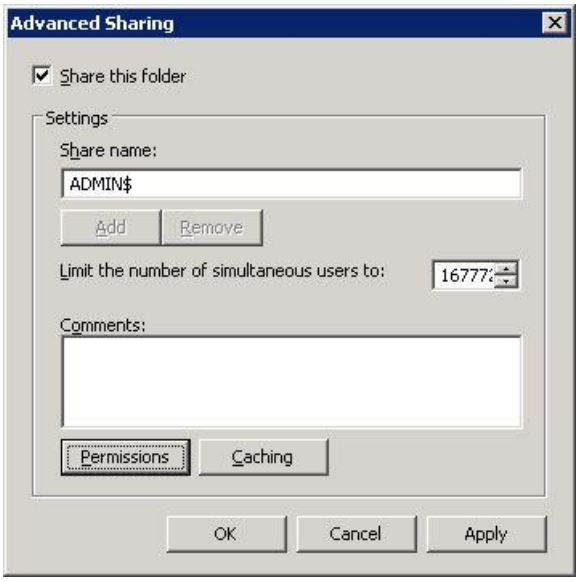

Figure 1.22: Enabling the ADMIN$ share

10. When Figure 1.23 appears, click on the **Add** button therein.

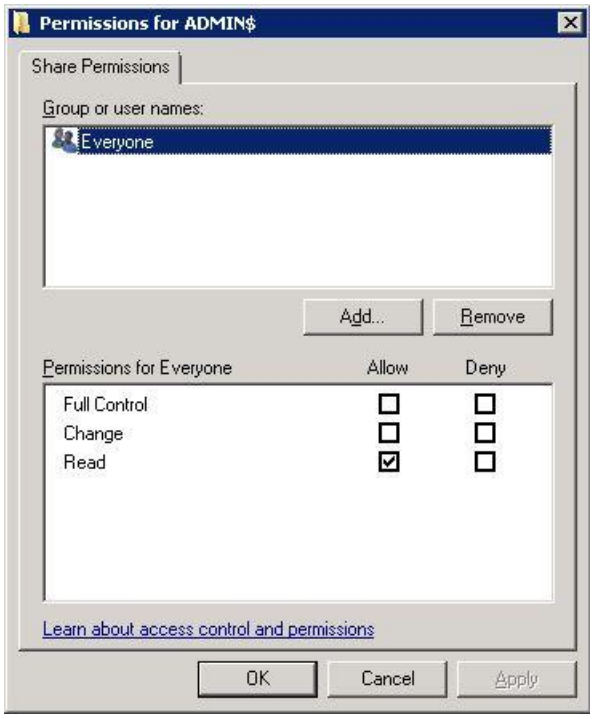

Figure 1.23: Clicking on the Add button

11. To allow a domain administrator to access the folder, first, ensure that a valid domain is specified in the **From this location** box of Figure 1.24. If you want to grant access to a local administrator instead, ensure that the name of the local host is displayed in the **From this location** box. To change this specification, use the **Locations** button in Figure 1.24. Then, enter the name of the local/domain administrator in the **Enter the object names to select** text area, and click the **OK** button.

Figure 1.24: Allowing a domain administrator to access the folder

12. The newly added user will be listed in the **Group or user names** section, as depicted by Figure 1.25. Select this user, and then, check all the three check boxes under **Allow** in the **Permissions for <user>** section in Figure 1.25. Then, click the **Apply** and **OK** buttons therein.

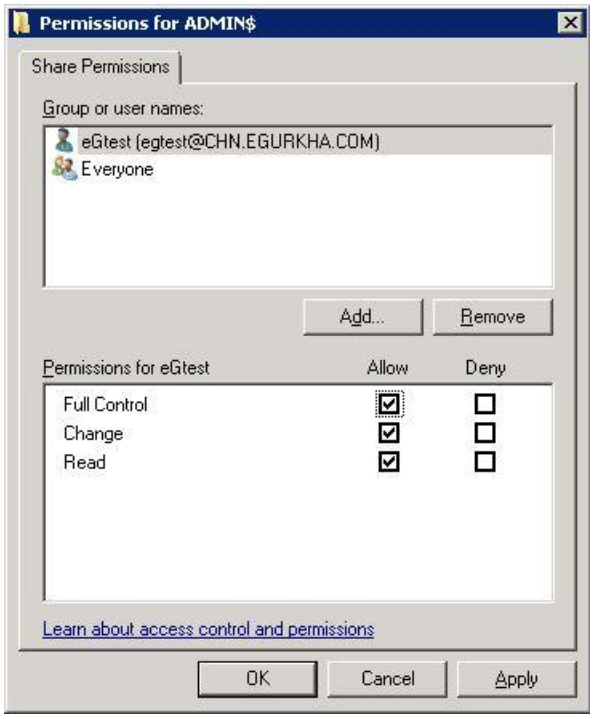

Figure 1.25: Allowing full access to the local/domain administrator

13. When Figure 1.26 appears, click on the **Apply** and **OK** buttons therein to register the changes.

Figure 1.26: Applying the changes

14. Alternatively, by adding a new entry in the Windows registry, you can quickly enable the **ADMIN$** share. The steps for the same are discussed hereunder:

    - In Run prompt type **regedit** to open registry editor.
    - Browse through the following sub key:

    HKEY_LOCAL_MACHINE\SOFTWARE\MICROSOFT\WINDOWS\CURRENTVERSION\POLICIES\SYSTEM

- Create a new entry with the below information
    - Key Name : LocalAccountTokenFilterPolicy
    - Key Type : DWORD (32-bit)
    - Key Value : 1
- Exit registry editor.

Note:

As with any change to the registry, ensure that the above-mentioned change is also performed with utmost care, so as to avoid problems in the functioning of the operating system.

Once the pre-requisites are fulfilled, you can proceed to use either of the monitoring models - *RHEV Hypervisor* or *RHEV Hypervisor - VDI* - to monitor the RHEV Hypervisor in your environment. The chapters that follow will discuss each of these models in detail.

Chapter 2

The RHEV Hypervisor Monitoring Model

As already mentioned, the *RHEV Hypervisor* model can be used where the VMs on the RHEV Hypervisor support critical server applications such as Web server, database server, mail server, etc.

Figure 2.1 depicts the monitoring model of the *RHEV Hypervisor*.

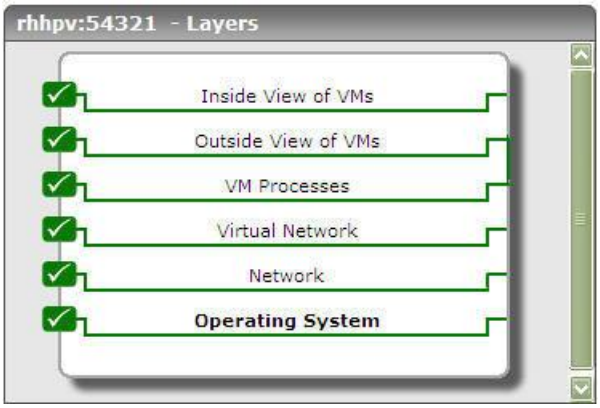

Figure 2.1: The layer model of the RHEV Hypervisor

Each layer of Figure 2.1 is mapped to a wide variety of tests that report a wealth of performance tests using which administrators can find quick and accurate answers to the following performance queries:

- What is the CPU load on the RHEV server and each of the VMs?
- What is the free physical memory in the RHEV server and which VM is contributing to the memory usage?
- Which network interfaces of the RHEV server are seeing the most traffic?
- How many VMs are running? What are their IP addresses/host names and operating systems?
- How many virtual CPUs are allocated to each VM?
- What portion of the physical server's CPU is used by each VM?
- How much memory is configured for each VM? What percentage of the configured memory is each VM consuming?
- How many disk reads and writes are being initiated by each VM?
- How much network traffic is being generated by every VM?

- What percentage of the physical CPU allocated to a VM is being used by the processes running on the VM?
- Which processes running in the VM are responsible for the resource usage (CPU, memory, disk) of the VM?
- Do all the disk partitions in the VM operating system have adequate space?
- Is there excessive queuing for disk access on any VM operating system? Which applications could be causing these excessive accesses?
- Are all critical Windows services running in the VM operating system?
- At what times of the day was the VM rebooted?
- Is any process running in the VM leaking memory or handles?

The section below will take a closer look at each layer of Figure 2.1.

2.1 The Operating System Layer

The tests mapped to this layer reveal the health of the RHEV hypervisor.

Figure 2.2: The tests mapped to the Operating System layer

2.1.1 Host Details - RHEV

This test proactively alerts administrators to the potential failure of an RHEV hypervisor by promptly capturing and reporting even the slightest change in the status of the hypervisor. In addition, the test also reports the number of physical CPUs the hypervisor has been configured with, and how the hypervisor's physical memory is being shared by the VMs - in the KSM mode or the THP mode?

Purpose	Proactively alerts administrators to the potential failure of an RHEV hypervisor by promptly capturing and reporting even the slightest change in the status of the hypervisor. In addition, the test also reports the number of physical CPUs the hypervisor has been configured with, and how the hypervisor's physical memory is being shared by the VMs - in the KSM mode or the THP mode?

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured
3. RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD - To auto-discover the VMs on a target RHEV hypervisor and obtain the outside view of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the outside view, you need to configure the test with the following:
➢ RHEL MGR HOST - The IP address/host name of the RHEV manager that the eG agent should connect to
➢ RHEL MGR PORT - The port number at which the said RHEV manager listens
➢ RHEL MGR DOMAIN - The domain to which the RHEV manager belongs
➢ RHEL MGR USER and RHEL MGR PASSWORD - The credentials of a user with read-only access to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document.
If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters.
If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the outside view, then, you can select any IP address of your choice from the RHEL MGR HOST list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an RHEL MGR HOST from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters.
On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the RHEL MGR HOST list, then, you can configure an RHEV manager on-the-fly by picking the Other option from the RHEL MGR HOST list. An ADD THE RHEV MANAGER DETAILS window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters.
4. CONFIRM PASSWORD - Confirm the RHEL MGR PASSWORD by retyping it here.
5. SSL - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to Yes. If not, set it to No.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the RHEV Hypervisor being monitored</td></tr>
<tr><td>Measurements made by the</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

<table>
<tr>
<td>test</td>
<td>Hypervisor status:
Indicates the current status of the RHEV hypervisor.</td>
<td></td>
<td>The values that this measure can report and their corresponding numeric values are as follows:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Up</td><td>1</td></tr>
<tr><td>Down</td><td>0</td></tr>
<tr><td>Error</td><td>2</td></tr>
<tr><td>Non responsive</td><td>3</td></tr>
<tr><td>Problematic</td><td>4</td></tr>
<tr><td>Not operational</td><td>5</td></tr>
<tr><td>Install failed</td><td>6</td></tr>
<tr><td>Installing</td><td>7</td></tr>
<tr><td>Reboot</td><td>8</td></tr>
<tr><td>Preparing for maintenance</td><td>9</td></tr>
<tr><td>Pending approval</td><td>10</td></tr>
<tr><td>Initializing</td><td>11</td></tr>
<tr><td>Maintenance</td><td>12</td></tr>
<tr><td>Unassigned</td><td>13</td></tr>
</table>
Note:
By default, this measure reports one of the Measure Values listed in the table above. The graph of this measure however will represent the hypervisor status using the numeric equivalents - ‘0’ t0 '13’.</td>
</tr>
<tr>
<td></td>
<td>Physical CPUs:
Indicates the number of physical CPUs available to the hypervisor.</td>
<td>Number</td>
<td></td>
</tr>
</table>

	Is this server a storage pool manager?: Indicates whether/not the host is currently the storage pool manager.		The Storage Pool Manager (SPM) coordinates all the metadata changes across the datacenter. This includes creating, deleting and manipulating virtual disks (Images), snapshots, and templates, and allocating storage for sparse block devices (on SAN). The SPM role is granted by the Red Hat Enterprise Virtualization Manager and can be migrated between any host in a data center. This means that all hosts in a data center must have access to all the storage domains defined in the data center. Red Hat Enterprise Virtualization Manager ensures that the SPM is always available and in case of errors will try to move the SPM role to a different host. This means that if the host that is running as the SPM has problems accessing the storage, the Manager will automatically check if there is another available host that can access the storage and will move the SPM over to that host. When the SPM starts, it tries to ensure that it is the only host that was granted the role, therefore it will acquire a storage-centric lease. This process can take some time. This measure reports the value *Yes* if the host is the storage pool manager. If the host experiences issues while accessing the storage domains, then the RHEV Manager will automatically move the SPM role to another available host. In this case, the value of this measure will be *No*. The numeric values that correspond to these measure values are as follows: **Measure Value** / **Numeric Value** Yes / 1 No / 0

			Note: By default, this measure reports one of the **Measure Value**s listed in the table above. The graph of this measure however will represent the SPM status using the numeric equivalents - '0' or '1'.
	Kernel samepage merging: Indicates whether/not **Kernel sampage merging** (KSM) is enabled for the RHEV hypervisor.		Memory page sharing is supported through a kernel feature called Kernel Same-page Merging (KSM). KSM scans the memory of each virtual machine and where virtual machines have identical memory pages KSM merges these into a single page that is shared between the virtual machines, storing only a single copy. If a guest attempts to change this shared page it will be given it's own private copy. When consolidating many virtual machines onto a host there are many situations in which memory pages may be shared – for example unused memory within a Windows virtual machine, common DLLs, libraries, kernels or other objects common between virtual machines. With KSM more virtual machines can be consolidated on each host, reducing hardware costs and improving server utilization. This measure reports the value *Enabled* if KSM is enabled on the hypervisor and reports *Disabled* if KSM is not enabled. The numeric values that correspond to these measure values are as follows: **Measure Value** / **Numeric Value** Enabled / 1 Disabled / 0

			Note: By default, this measure reports one of the **Measure Value**s listed in the table above. The graph of this measure however will represent the KSM status using the numeric equivalents - '0' or '1'.
	Transparent hugepages: Indicates whether the Transparent hugepages (THP) memory management mechanism has been enabled on the hypervisor or not .		Typically, there are two ways to enable the system to manage large amounts of memory: ➢ Increase the number of page table entries in the hardware memory management unit ➢ Increase the page size Since the first method is expensive, RHEL implements the second method via *huge pages*. Huge pages are blocks of memory that come in 2MB and 1GB sizes. The page tables used by the 2MB pages are suitable for managing multiple gigabytes of memory, whereas the page tables of 1GB pages are best for scaling to terabytes of memory. Huge pages must be assigned at boot time. They are also difficult to manage manually, and often require significant changes to code in order to be used effectively. As such, Red Hat Enterprise Linux 6 also implemented the use of *transparent huge pages* (THP). THP is an abstraction layer that automates most aspects of creating, managing, and using huge pages. THP hides much of the complexity in using huge pages from system administrators and developers. The default settings of THP improves the performance of most system configurations. Note that THP can currently only map anonymous memory regions such as heap and stack space.

			If THP is enabled on the hypervisor, then this measure reports the value *Enabled*. If not, the value of this measure will be *Disabled*. The numeric values that correspond to these measure values are as follows: **Measure Value** / **Numeric Value** Enabled / 1 Disabled / 0 **Note:** By default, this measure reports one of the **Measure Value**s listed in the table above. The graph of this measure however will represent the THP status using the numeric equivalents - '0' or '1'.

2.1.1.1 Configuring an RHEV Manager to Use for Monitoring the RHEV Hypervisor

To configure an RHEV manager on-the-fly for use to monitor an RHEV Hypervisor, do the following:

1. From the **RHEL MGR HOST** parameter in the test configuration page, select the **Other** option.
2. Figure 2.3 will then appear:

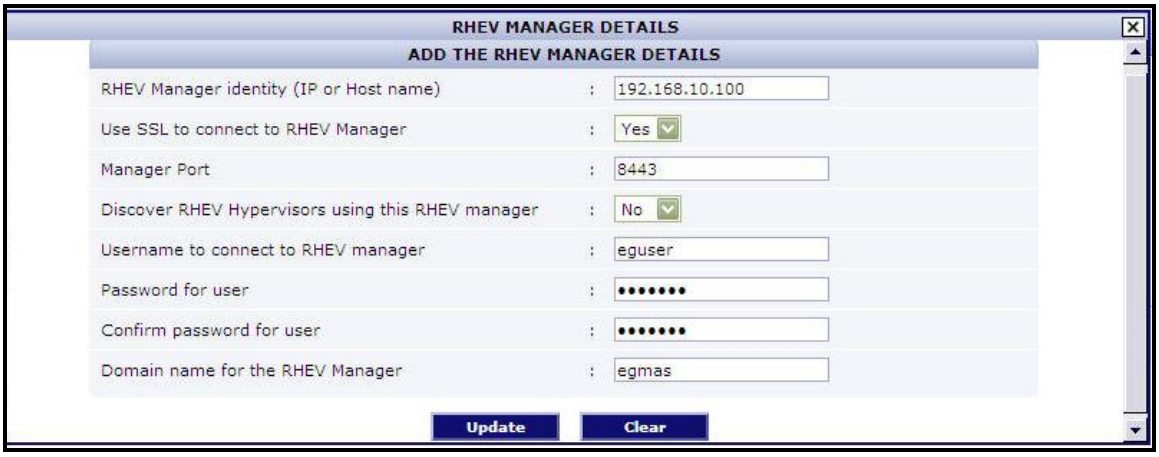

Figure 2.3: Configuring the details of the RHEV Manager

3. Specify the following in Figure 2.3:

- **RHEV Manager Identify**: Specify the IP address/host name of the RHEV manager in your environment.
- **Use SSL to Connect to the RHEV Manager**: Set this flag to **Yes** if the RHEV manager in your environment is SSL-enabled. Otherwise, set this flag to **No**.
- **Manager Port**: If the RHEV manager is SSL-enabled, then 8443 will be displayed here by default. On the other hand, if the manager is not SSL-enabled, the default **Manager Port** will be 8080. If the RHEV manager in your environment listens on a different SSL or non-SSL port, then make corresponding changes to the default setting.
- **Discover RHEV Hypervisors using this RHEV Manager**: If you also want to discover additional RHEV servers in your environment using this RHEV manager, set this flag to **Yes**. If you only want to use this RHEV manager to obtain the *outside view* of VMs, set this flag to **No**.
- **Username to connect to RHEV Manager** and **Password for user**: Specify the credentials (i.e., user name and password) of a user who has been assigned *read-only* access. To create a *read-only* role and assign it to a user, follow the steps detailed in Section 1.4 of this document.
- **Confirm password for user**: Confirm the password of the **RHEVMUser** by retyping it here.
- **Domain name for the RHEV manager**: Specify the name of the domain to which the RHEV manager belongs.

4. Once the details required by Figure 2.3 are provided, click the **Update** button therein to proceed.
5. This will take you back to the test configuration page. However, this time, you will find that the **RHEL MGR HOST**, **RHEL MGR PORT**, **RHEL MGR USER**, **RHEL MGR PASSWORD**, **RHEL MGR DOMAIN**, and **SSL** parameters in the page are all configured with the values passed to the corresponding fields in Figure 2.3.

2.1.2 Memory Details - RHEV Test

A contention for memory resources on the RHEV hypervisor can affect the memory allocation to VMs, which in turn can adversely impact the performance of the applications operating on the VMs. It is therefore imperative that you closely observe how the hypervisor uses the physical memory available to it, so that you can proactively determine a potential memory crunch. This can be achieved with the help of the **Memory Details - RHEV** test. This test periodically monitors the memory usage of the hypervisor, checks whether adequate free memory is available to the hypervisor, and if not, promptly alerts users to the same. In the process, the test also reveals the top memory VMs executing on the hypervisor, checks swap memory usage, and also reports whether/not memory has been overcommitted by the hypervisor. This way, the test also points you to the probable reasons for the memory erosion (if any) - is it owing to memory-starved VMs on the hypervisor? or is it because the hypervisor has overcommitted memory?

Purpose	Periodically monitors the memory usage of the hypervisor, checks whether adequate free memory is available to the hypervisor, and if not, promptly alerts users to the same

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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. The RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

	6. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG 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 the RHEV hypervisor being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Total physical memory: Indicates the total physical memory of the RHEV hypervisor.	GB	
	Free physical memory: Indicates the amount of physical memory currently unused on the hypervisor.	GB	Ideally, the value of this measure should be high. A low value or a consistent decrease in this value could indicate a gradual memory erosion, which can consequently affect the host as well as VM performance.
	Used physical memory: Indicates the amount of physical memory currently being utilized by the hypervisor.	GB	Ideally, the value of this measure should be low. A very high value or a value that grows dangerously close to the *Total physical memory* of the host is a cause for concern.
	Memory overhead: Indicates the memory overhead on the hypervisor.	GB	The value of this measure is typically the difference between the *Total physical memory* and the sum of the *Used physical memory* and the *Free physical memory* measures.

	Physical memory usage: Indicates the percentage of total physical memory that is being used by the hypervisor.	Percent	Ideally, the value of this measure should be low. A high value or a value close to 100% indicates a contention for memory resources on the hypervisor. Use the detailed diagnosis of this measure to know the memory configuration of each VM on the RHEV hypervisor, and the how every VM is using the configured memory. Memory-hungry VMs can thus be isolated.
	Physical memory free: Indicates the percentage of percentage of total physical memory that is currently free on the hypervisor.	Percent	Ideally, the value of this measure should be high. A low value or a consistent decrease in this value indicates a contention for memory resources on the hypervisor.
	I/O memory buffers: Indicates the total I/O memory buffers in the hypervisor.	MB	
	Cached memory: Indicates the total OS cached memory in the RHEV hypervisor.	MB	
	Total swap memory: Indicates the total amount of swap memory on the hypervisor.	GB	
	Free swap memory: Indicates the amount of swap memory that is currently unused on the hypervisor.	GB	A high value is desired for this measure, as a consistent decrease in this value is a sign of excessive swap usage, which in turn signals a memory bottleneck.
	Used swap memory: Indicates the amount of swap memory that is currently utilized by the hypervisor.	GB	Significant or consistent memory swapping indicates that the hypervisor is severely overcommitted and that performance degradation is imminent or actively occurring.
	Cached swap memory: Indicates the total swap memory that the hypervisor has cached.	MB	

<table>
<tr><td></td><td>Swap memory usage:
Indicates the percentage of swap memory used by the hypervisor.</td><td>Percent</td><td>A high value is indicative of a contention for memory resources on the hypervisor.</td></tr>
<tr><td></td><td>Is memory overcommitted?:
Indicates whether the hypervisor memory is overcommitted or not.</td><td></td><td>Hypervisor memory is over-committed when the total memory space allocated (memory granted) to powered-on VMs, plus hypervisor memory overhead, is greater than the amount of total physical memory available to the host. A severe memory overcommitment can cause serious performance degradations.

If the value of this measure is Yes, it indicates that the hypervisor memory is overcommitted. The value No indicates that there is no overcommitment of memory.

The numeric values that correspond to these measure values are indicated in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

By default, this measure reports one of the Measure Values listed in the table above. The graph of this measure however will represent the memory overcommitment status using the numeric equivalents - ‘0’ or '1’.</td></tr>
</table>

2.1.3 CPU Details - RHEV Test

Excessive usage of physical CPU resources by the RHEV hypervisor or its VMs can be the source of prolonged slowdowns that may be experienced by those VMs. Hence, whenever users complaint of slowdowns in their virtual applications, it would be best to first check whether the hypervisor has adequate unused CPU resources, as a CPU contention can have a disastrous effect on VM performance and consequently application performance. This is why, the eG agent, with the help of its **CPU Details - RHEV** test, runs periodic usage checks on the CPU resources of the hypervisor. Besides proactively detecting abnormal CPU consumption by the hypervisor, the test also accurately points you to the root-cause of the CPU contention - did it happen because of CPU-hungry VMs on the hypervisor? did it happen because of CPU-hungry user processes or system-level processes? did it occur when the hypervisor performed Kernel Same-Page Merging? or did it happen when the CPU was idle?

Purpose	Runs periodic usage checks on the CPU resources of the hypervisor. Besides proactively detecting abnormal CPU consumption by the hypervisor, the test also accurately points you to the root-cause of the CPU contention - did it happen because of CPU-hungry VMs on the hypervisor? did it happen because of CPU-hungry user processes or system-level processes? did it occur when the hypervisor performed Kernel Same-Page Merging? or did it happen when the CPU was idle?
Configurable parameters for the test	1. **NTEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **RHEL MGR HOST**, **RHEL MGR PORT**, **RHEL MGR DOMAIN**, **RHEL MGR USER**, **RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

	6. **HYPERVISOR USER** - Specify the name of a user who has the right to connect to the RHEV hypervisor via SSH. 7. **HYPERVISOR PASSWORD** - Specify the password of the **HYPERVISOR USER**. 8. **CONFIRM PASSWORD** - Confirm the **HYPERVISOR PASSWORD** by retyping it here. 9. **HYPERVISOR SSH PORT** - Enter the SSH port at which the RHEV hypervisor listens. 10. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG 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 the RHEV hypervisor being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	User CPU utilization: Indicates the percentage of CPU utilized by user processes.	Percent	Comparing the value of these measures will enable administrators to figure out where the hypervisor is spending the maximum CPU time - in processing user requests? in system-level processes? or when being idle?
	System CPU utilization: Indicates the percentage of CPU resources that the hypervisor utilized for system-level processing.	Percent	
	Idle CPU utilization: Indicates the percentage of CPU time utilized when the hypervisor was idle.	Percent	

	CPU utilization: Indicates the percentage of CPU utilized by the hypervisor.	Percent	A high value of this measure is a cause for concern, as it indicates excessive CPU usage by the hypervisor. If left unchecked, it may cause a serious contention for CPU resources amidst VMs. Use the detailed diagnosis of this measure to know which VMs on the hypervisor are consuming the CPU resources excessively.
	Kernel samepage merging CPU utilization: Indicates the percentage of CPU time spent by the hypervisor when performing Kernel Same-page Merging (KSM).	Percent	Memory page sharing is supported through a kernel feature called Kernel Same-page Merging (KSM). KSM scans the memory of each virtual machine and where virtual machines have identical memory pages KSM merges these into a single page that is shared between the virtual machines, storing only a single copy. If a guest attempts to change this shared page it will be given it's own private copy. When consolidating many virtual machines onto a host there are many situations in which memory pages may be shared – for example unused memory within a Windows virtual machine, common DLLs, libraries, kernels or other objects common between virtual machines. With KSM more virtual machines can be consolidated on each host, reducing hardware costs and improving server utilization. Use this measure to determine how CPU-intensive KSM is.

2.2 The Network Layer

Using the tests mapped to this layer, know whether the RHEV hypervisor is available over the network or not, and promptly detect network interfaces that are down, slow, or are using bandwidth excessively. Since the *Network* test mapped to this layer has already been discussed in the *Monitoring Unix and Windows Servers* document, let us take a look at the *Network - RHEV* test alone.

Figure 2.4: The tests mapped to the Network layer

2.2.1 Network - RHEV Test

This test auto-discovers the network interfaces supposed by the RHEV server and reports the current status and speed of every discovered interface, and the errors encountered by each. This way, the test sheds light on the slow, error-prone, and congested (in terms of level of network traffic) network interfaces on the hypervisor.

Purpose	Auto-discovers the network interfaces supposed by the RHEV server and reports the current status and speed of every discovered interface, and the errors encountered by each.

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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

<table>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every network interface supported by the RHEV hypervisor being monitored</td></tr>
<tr><td rowspan="4">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Status:
Indicates the current status of this network interface.</td><td></td><td>If the network interface is up and running, then the value of this measure will be Up. On the other hand, if the network interface is currently non-operational, then this measure will report the value Down.
The numeric values that correspond to the measure values mentioned above are as follows:<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Up</td><td>1</td></tr><tr><td>Down</td><td>0</td></tr></table>Note:
By default, this measure reports one of the Measure Values listed in the table above. The graph of this measure however will represent network interface status using the numeric equivalents - ‘0’ or '1’.</td></tr>
<tr><td>Speed:
Indicates the current speed of this network interface.</td><td>Mbps</td><td>Compare the value of this measure across interfaces to determine the slowest interface.</td></tr>
<tr><td>Network data transmitted:
Indicates the rate at which data is transmitted over this network interface.</td><td>Mbps</td><td rowspan="2">A high rate of incoming and outgoing data could indicate that the network interface is experiencing high levels of network traffic.</td></tr>
<tr><td>Network data received:
Indicates the rate at which data is received over this network interface.</td><td>Mbps</td></tr>
</table>

	Errors during transmission: Indicates the number of errors that occurred when data was transmitted over this interface.	Number	Ideally, the value of both these measures should be 0. Comparing the value of each of these measures across interfaces will introduce you to the error-prone interfaces.
	Errors during reception: Indicates the number of errors that occurred when data was received over this interface.	Number	

2.3 The Virtual Network Layer

The test mapped to this layer reveals the level of network traffic transacted over the virtual networks and errors experienced by these networks.

Figure 2.5: The tests mapped to the Virtual Network layer

2.3.1 RHEV Virtual Network Traffic Test

Red Hat Enterprise Virtualization Manager supports multiple virtual networks and VLANs, allowing an administrator to centrally manage and configure the virtual network. Virtual networks emulate network connectivity within the RHEV server and allow VMs hosted on that server to exchange data. Continuous monitoring of these virtual networks will enable administrators to isolate which virtual networks are healthy and are too busy trafficking data to and from VMs, and which virtual networks are experiencing error conditions. The **Virtual Network Traffic** test does just that.

Purpose	Helps isolate which virtual networks are healthy and are too busy trafficking data to and from VMs, and which virtual networks are experiencing error conditions

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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

Outputs of the test	One set of results for every virtual network on the RHEV hypervisor being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Network data transmitted: Indicates the rate at which data is transmitted over this virtual network.	Mbps	A high rate of incoming and outgoing data could indicate that the virtual network is very busy or is been overloaded with traffic.
	Network data received: Indicates the rate at which data is received over this virtual network.	Mbps	
	Errors during transmission: Indicates the number of errors that occurred when data was transmitted over this virtual network.	Number	Ideally, the value of both these measures should be 0. Comparing the value of each of these measures across virtual networks will introduce you to the error-prone networks.
	Errors during reception: Indicates the number of errors that occurred when data was received over thisbvirtual network.	Number	

2.4 The VM Processes Layer

The test mapped to this layer monitors the availability and responsiveness of configured ports on the RHEV hypervisor. Since this test has already been discussed in the *Monitoring Unix and Windows Servers* document, let us switch to the next layer.

Figure 2.6: The test mapped to the VM Processes layer

2.5 The Outside View of VMs Layer

This layer provides the host operating system's view of the resource usage levels of each of the virtual guests hosted on it. Using the information reported by this test, administrators can:

- Determine which of the guests is taking up more resources (CPU, memory, network, or disk) than the others. This information can help with load balancing or capacity planning. For example, if one of the guests is receiving a very high rate of requests compared to the others, this guest may be a candidate for migration to another RHEV server, so as to minimize the impact it has on the other guests on the current RHEV server.
- Determine times when sudden or steady spikes in the physical resource utilization are caused by the guest machines
- Track the overall status of the virtual machines - how many are registered, which ones are powered on, and at what times, etc.

Figure 2.7: The tests mapped to the Outside View of VMs layer

2.5.1 RHEV VM Details Test

This test monitors the amount of the physical server's resources that each guest on an RHEV server is taking up. Using the metrics reported by this test, administrators can determine which virtual guest is taking up most CPU, which guest is generating the most network traffic, which guest is over-utilizing memory, etc. Note that the amount of resources taken up by a virtual guest will be limited by the resource allocations that have been made by administrators. For example, an administrator could cap the amount of memory that a specific guest may take.

Purpose	Monitors the amount of the physical server's resources that each guest on an RHEV server is taking up

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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

<table>
<tr><td></td><td colspan="3">6. HYPERVISOR USER - Specify the name of a user who has the right to connect to the RHEV hypervisor via SSH.
7. HYPERVISOR PASSWORD - Specify the password of the HYPERVISOR USER.
8. CONFIRM PASSWORD - Confirm the HYPERVISOR PASSWORD by retyping it here.
9. HYPERVISOR SSH PORT - Enter the SSH port at which the RHEV hypervisor listens.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every VM on the RHEV hypervisor being monitored</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>VM power state:
Indicates whether this VM is currently powered-on or off.</td><td></td><td>This measure reports the value Up if the VM is currently powered-on, and the value Down if the VM is currently powered-off.
The numeric values that correspond to the measure values mentioned above are as follows:<table><tr><th>Numeric Value</th><th>Measure Value</th></tr><tr><td>1</td><td>Up</td></tr><tr><td>0</td><td>Down</td></tr></table>Note:
By default, this measure reports the above-mentioned Measure Values while indicating the status of this VM. However, the graph of this measure powered-on states will be represented using the corresponding numeric equivalents only.</td></tr>
</table>

	Is stateless VM?: Indicates whether this VM is currently stateless or not.		A stateless VM is not a VM that has its own local data. More often than not, it does require some local data or a local cache for better performance. But these data don't need to be persisted. In some cases, a stateless VM can have additional software installed or data pulled in from a known repository. This process should be fully automated with self-starter scripts, or managed by an external installer. Once a stateless VM goes live, it should discover all the related services to persistent data. The stateless VM has to rely on the environment to work effectively. It includes directory services, data services, and so on. With stateless VMs, you can improve mobility inside an enterprise and external transfer to the public cloud. For one thing, you just need to transfer a VM image once and only once. When your application runs into problems, instead of diagnosing the problem you just remove the problematic VMs and add new virtual machines. With this capability, you can also easily scale out your applications by adding new VM instances as you need them. Last but not least, the software upgrade and patch. It has been a big pain to upgrade and patch software in large deployments. You have to do it with each individual machine despite virtual or not. With stateless VM, you only need to patch the template and new virtual machines will pick it up seamlessly. This measure reports the value *Yes* if the VM is a stateless VM, and the value *No* if it is not a stateless VM. The numeric values that correspond to these measure values are discussed in the table below:

			<table><tr><th>Numeric Value</th><th>Measure Value</th></tr><tr><td>1</td><td>Yes</td></tr><tr><td>0</td><td>No</td></tr></table> **Note:** By default, this measure reports the above-mentioned **Measure Values** while indicating the VM is a stateless VM or not. However, in the graph of this measure this will be represented using the corresponding numeric equivalents only.
	Number of VCPU: Indicates the number of virtual CPUs allocated to this VM.	Number	
	System usage of physical CPU: Indicates the percentage of physical CPU resources this VM utilized for system-level processing.	Percent	A high value could indicate that the VM is executing too many system-level tasks simultaneously.
	Virtual CPU utilization: Indicates the percentage of virtual CPU resources this VM utilized.	Percent	Compare the value of this measure across VMs to identify the VM that is consuming CPU excessively. A high value for this measure could indicate that one/more CPU-intensive processes are executing on the VM.
	Configured memory: Indicates the amount of memory that is allocated to this VM.	MB	
	Physical memory consumed: Indicates the amount of physical memory consumed by this VM.	MB	

<table>
<tr><td></td><td>Memory usage:
Indicates the percentage of physical memory consumed by this VM.</td><td>Percent</td><td>A high value for this measure is indicative of high memory usage by a VM. Compare the value of this measure across VMs to know which VMs are eroding the physical memory of the hypervisor. Once the resource-hungry VMs are isolated, you need to investigate why those VMs are consuming memory excessively and see how the resource usage can be controlled. If the issue is allowed to persist, then very soon you may not have adequate physical memory to support hypervisor and VM operations.</td></tr>
<tr><td></td><td>Guaranteed memory:
Indicates the amount of memory resources that is guaranteed available to this VM - i.e., the minimum amount of memory that will always be available to this VM.</td><td>MB</td><td></td></tr>
<tr><td></td><td>Disk status:
Indicates the status of the virtual disk of this VM.</td><td></td><td>The value of this measure can be OK or Not ok, depending upon the current status of the disk.
The numeric values that correspond to these measure values are as follows:<table><tr><th>Numeric Value</th><th>Measure Value</th></tr><tr><td>1</td><td>Ok</td></tr><tr><td>0</td><td>Not ok</td></tr></table>Note:
By default, this measure reports the Measure Values listed in the table above to indicate the status of the disk. However, in the graph of this measure, the disk status will be represented using the numeric equivalents of the measure values only.</td></tr>
<tr><td></td><td>Disk capacity:
Indicates the current disk capacity of this VM.</td><td>GB</td><td></td></tr>
</table>

	Data reads from disk: Indicates the rate at which data is read from the virtual disks of this VM.	MB/Sec	
	Data writes to disk: Indicates the rate at which data is written to the virtual disks of this VM.	MB/Sec	
	Disk throughput: Indicates the rate at which I/O operations are performed on the virtual disks of this VM.	MB/Sec	The value of this measure indicates the level of I/O activity on every VM. Compare this value across VMs to identify which VM is experiencing abnormally high disk I/O. Zooming into the internal operations of that VM can shed light on the I/O-intensive processes that may be executing in that VM.
	Network data transmitted: Indicates the rate at which data is transmitted from this VM.	Mbps	
	Network data received: Indicates the rate at which data is received by this VM.	Mbps	
	Network throughput: Indicates the rate at which network data is accessed by this VM.	Mbps	For every VM, the value of this measure indicates the level of network traffic flowing into and from that VM. Compare this value across VMs to identify which VM is experiencing abnormally high traffic.

2.5.2 RHEV VM Status Test

The value of this measure indicates the level of I/O activity on every VM. Compare this value across VMs to identify which VM is experiencing abnormally high disk I/O. Zooming into the internal operations of that VM can shed light on the I/O-intensive processes that may be executing in that VM.

Purpose	Monitors the amount of the physical server's resources that each guest on an RHEV server is taking up

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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

	6. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG 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 the RHEV hypervisor being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Registered VMs: Indicates the total number of VMs that have been registered with the RHEV hypervisor.	Number	Use the detailed diagnosis of this measure to view the details of the registered VMs.
	Powered on VMs: Indicates the number of VMs that are currently powered on.	Number	To know which are the VMs that are powered on, use the detailed diagnosis capability of this measure (if enabled).
	Powered off VMs: Indicates the number of VMs that are currently powered off.	Number	To know which are the VMs that are powered on, use the detailed diagnosis capability of this measure (if enabled).
	Suspended VMs: Indicates the number of VMs that are currently in a **suspended** state.	Number	Use the detailed diagnosis of this measure to know which VMs are currently in a **suspended** state. The suspend action saves the virtual machine state to disk and stops it.
	Orphaned VMs: Indicates the number of VMs that are currently in an UNKNOWN state.	Number	Use the detailed diagnosis of this measure to know which VMs are currently in an **unknown** state.
	Other VMs: Indicates the number of VMs currently in the **image_illegal** and **not_responding** states,.	Number	Use the detailed diagnosis of this measure to know which VMs are currently in an **image_illegal** and **not_responding** state.

	VM templates: Indicates the number of template VMs currently on the hypervisor.	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. You can use the detailed diagnosis of this measure to view the names and IP addresses of the template VMs.
	Added VMs: Indicates the number of VMs that were newly added to the hypervisor since the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which VMs were recently migrated to the hypervisor.
	Removed VMs: Indicates the number of VMs that were newly removed from the hypervisor since the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which VMs were recently migrated from the hypervisor.
	VMs with users assigned: Indicates the number of VMs that have been assigned to users.	Number	To know which VMs have been assigned to users, use the detailed diagnosis capability of this measure (if enabled).
	VMs with users unassigned: Indicates the number of VMs that have not been assigned to users.	Number	To know which VMs have not been assigned to users, use the detailed diagnosis capability of this measure (if enabled).

The detailed diagnosis of the *Registered VMs* measure reveals the name, IP address, and operating system of the registered VMs.

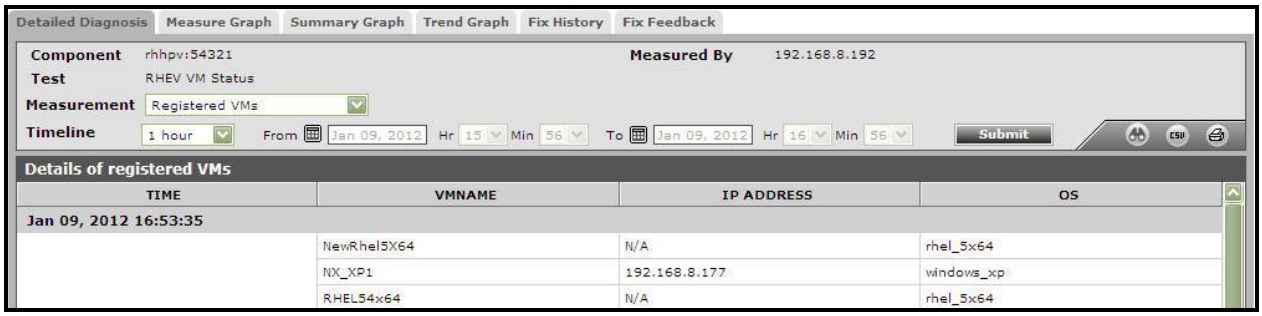

Figure 2.8: The detailed diagnosis of the Registered VMs measure

The detailed diagnosis of the *Orphaned VMs* measure reveals the name, IP address, and operating system of the VMs that are currently in an **Unknown** state.

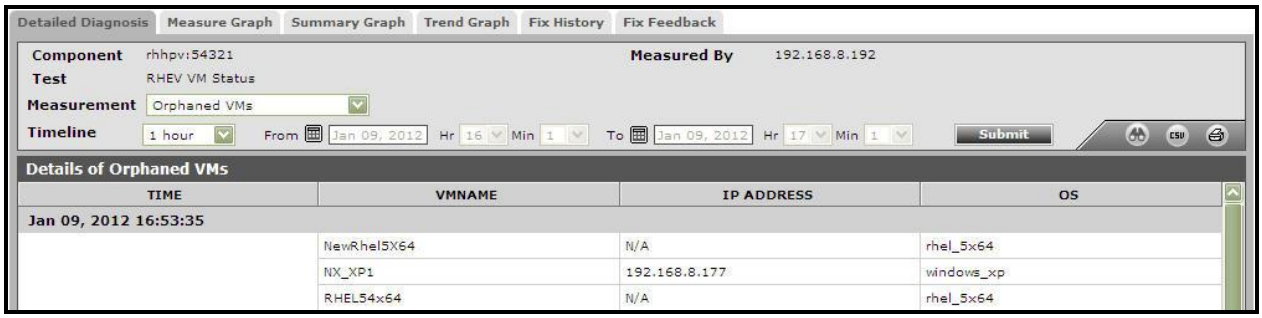

Figure 2.9: The detailed diagnosis of the Orphaned VMs measure

The detailed diagnosis of the *Powered on VMs* measure reveals the name, IP address, and operating system of the VMs that are currently powered on.

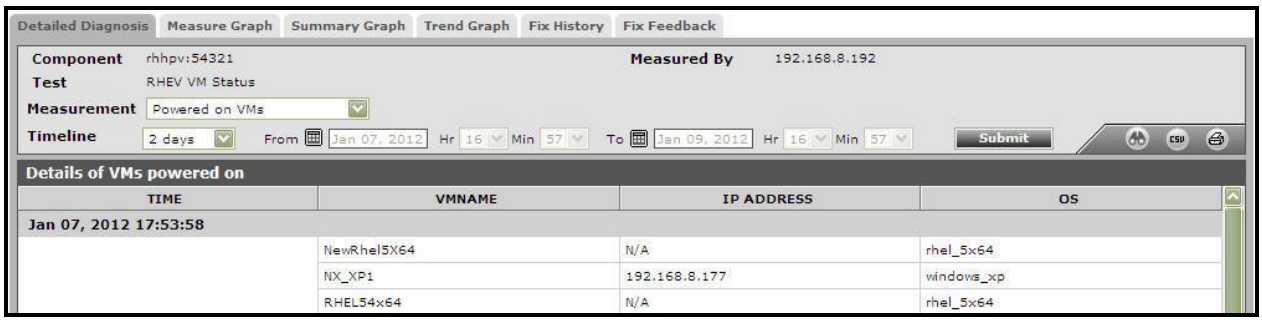

Figure 2.10: The detailed diagnosis of the Powered on VMs measure

The detailed diagnosis of the *VMs without users assigned* measure reveals the name, IP address, and operating system of the VMs that have been not been assigned to any users.

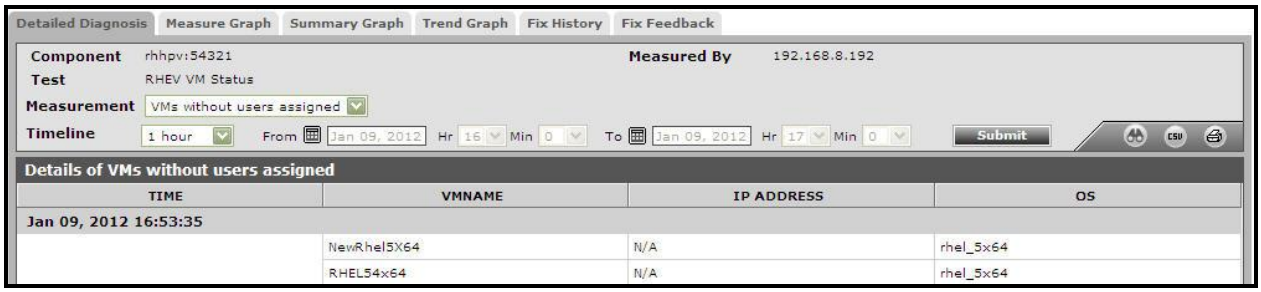

Figure 2.11: The detailed diagnosis of the Vms without users assigned measure

2.6 The Inside View of VMs Layer

The **Outside View of VMs** layer provides an "external" view of the different VM guests - the metrics reported at this layer are based on what the RHEV Hypervisor seeing about the performance of the individual guests. However, an external view of the VM guest operating system and its applications may not be sufficient. For instance, suppose one of the disk partitions of the guest operating system has reached capacity. This information cannot be gleaned from host operating system. Likewise, bottlenecks such as a longer process run queue or a higher disk queue length are more visible using an internal monitor. Internal monitoring (from within the guest operating system) also provides details about the resource utilization of different application(s) or processes.

The tests mapped to the **Inside View of VMs** layer provide an "internal" view of the workings of each of the guests - these tests execute on an RHEV hypervisor but send probes into each of the guest operating systems to analyze how well each guest utilizes the resources that are allocated to it, and how well it handles network traffic and loading.

By default however, clicking on the **Inside View of VMs** layer, does not display the list of tests associated with that layer. Instead, Figure 2.12 appears, which provides you with an overview of individual guest performance.

Figure 2.12: A list of guest operating systems on an RHEV server and their current state

To return to the layer model of the *RHEV Hypervisor* and view the tests associated with the **Inside View of VMs** layer, click on the **COMPONENT LAYERS** link in Figure 2.12. You can now view the list of tests mapped to the **Inside View of VMs** layer, as depicted by Figure 2.13 below.

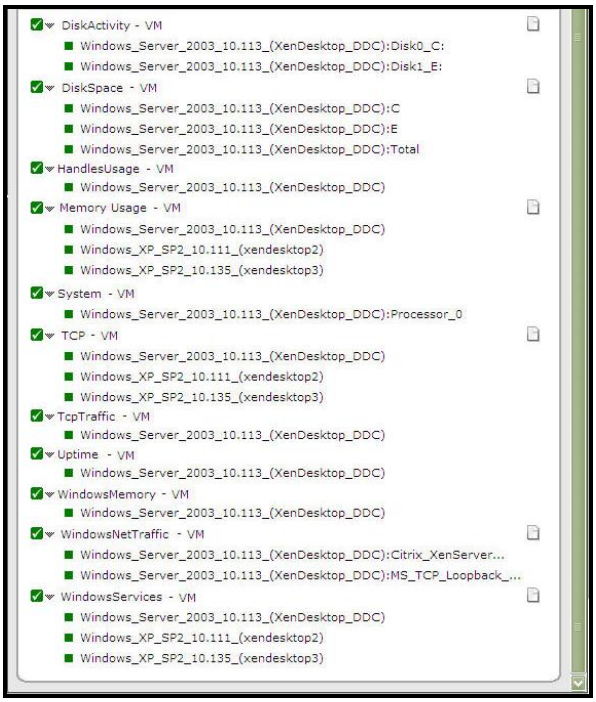

Figure 2.13: The tests associated with the Inside View of VMs layer

As indicated in Figure 2.13, the tests associated with this layer monitor different aspects of each virtual guest. Disk space utilization, disk activity levels, CPU utilization, memory usage levels, network traffic, etc. are all monitored and reported for each virtual guest hosted on the RHEV hypervisor. Detailed diagnosis for these tests provide details of individual processes and their utilization levels.

The tests associated with this layer are described in detail below.

2.6.1 Disk Activity - VM Test

This test reports statistics pertaining to the input/output utilization of each physical disk on a guest.

Purpose	To measure the input/output utilization of each physical disk on each guest of an RHEV Hypervisor
Target of the test	An RHEV Hypervisor
Agent deploying the test	A remote 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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

6. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

7. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

   Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights**. Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

8. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

   - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
   - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

     Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

     If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

<table>
<tr><td></td><td colspan="3">

If the guests belong to different domains - In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple DOMAIN names, multiple ADMIN USER names and ADMIN PASSWORDS would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page.

To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document.
If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' - In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to none.

9. REPORT BY USER - While monitoring a RHEV Hypervisor, the REPORT BY USER flag is set to NO by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a RHEV Hypervisor - VDI, this flag is set to YES by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every username_on_virtualmachinename.
10. REPORT POWERED OS - This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.

If the REPORT POWERED OS flag is set to Yes (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their virtual machine name and not by the username_on_virtualmachinename. On the other hand, if the REPORT POWERED OS flag is set to No, then this test will not report measures for those VMs to which no users are logged in currently.
11. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG 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.

</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every combination of virtual_guest:disk_partition</td></tr>
<tr><td>Measurements made by the</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

test	**Percent virtual disk busy:** Indicates the percentage of elapsed time during which the disk is busy processing requests (i.e., reads or writes).	Percent	Comparing the percentage of time that the different disks are busy, an administrator can determine whether load is properly balanced across the different disks.
	Percent reads from virtual disk: Indicates the percentage of elapsed time that the selected disk drive is busy servicing read requests.	Percent	
	Percent writes to virtual disk: Indicates the percentage of elapsed time that the selected disk drive is busy servicing write requests.	Percent	
	Virtual disk read time: Indicates the average time in seconds of a read of data from the disk.	Secs	
	Virtual disk write time: Indicates the average time in seconds of a write of data from the disk.	Secs	
	Avg. queue for virtual disk: Indicates the average number of both read and write requests that were queued for the selected disk during the sample interval.	Number	

	Current queue for virtual disk: The number of requests outstanding on the disk at the time the performance data is collected.	Number	This measure includes requests in service at the time of the snapshot. This is an instantaneous length, not an average over the time interval. Multi-spindle disk devices can have multiple requests active at one time, but other concurrent requests are awaiting service. This counter might reflect a transitory high or low queue length, but if there is a sustained load on the disk drive, it is likely that this will be consistently high. Requests experience delays proportional to the length of this queue minus the number of spindles on the disks. This difference should average less than two for good performance.
	Reads from virtual disk: Indicates the number of reads happening on a logical disk per second.	Reads/Sec	A dramatic increase in this value may be indicative of an I/O bottleneck on the guest.
	Data reads from virtual disk: Indicates the rate at which bytes are transferred from the disk during read operations.	KB/Sec	A very high value indicates an I/O bottleneck on the guest.
	Writes to virtual disk: Indicates the number of writes happening on a local disk per second.	Writes/Sec	A dramatic increase in this value may be indicative of an I/O bottleneck on the guest.
	Data writes to virtual disk: Indicates the rate at which bytes are transferred from the disk during write operations.	KB/Sec	A very high value indicates an I/O bottleneck on the guest.
	Disk service time: Indicates the average time that this disk took to service each transfer request (i.e., the average I/O operation time)	Secs	A sudden rise in the value of this measure can be attributed to a large amount of information being input or output. A consistent increase however, could indicate an I/O processing bottleneck.

	Disk queue time: Indicates the average time that transfer requests waited idly on queue for this disk.	Secs	Ideally, the value of this measure should be low.
	Disk I/O time: Indicates the average time taken for read and write operations of this disk.	Secs	The value of this measure is the sum of the values of the Disk service time and Disk queue time measures. A consistent increase in the value of this measure could indicate a latency in I/O processing.

The detailed diagnosis of the *Percent virtual disk busy* measure, if enabled, provides information such as the Process IDs executing on the disk, the Process names, the rate at which I/O read and write requests were issued by each of the processes , and the rate at which data was read from and written into the disk by each of the processes. In the event of excessive disk activity, the details provided in the detailed diagnosis page will enable users to figure out which process is performing the I/O operation that is keeping the disk busy. **The detailed diagnosis for this test is available for Windows guests only, and not Linux guests**.

Shows the IO operations done by the processes

Time	ID_Process	ProcessName	IO_Rate (Bytes/sec)	IO_Read_Rate (Bytes/sec)	IO_Read_Ops_Rate (Ops/Sec)	IO_Write_Rate (Bytes/sec)	IO_Write_Ops_Ra (Ops/sec)
Jan 03, 2008 05:44:08							
	696	services	30108.21	252.2	3.34	29856.01	12.68
	4	System	28489.58	28489.58	47.71	0	0
	1032	svchost	16500.6	15801.37	21.02	699.23	3
	628	csrss	3320	3320	7.01	0	0
	2396	vmiprvse	194.82	125.43	1.33	69.39	1.33
Jan 03, 2008 05:33:43							
	1032	svchost	770.92	137.38	3	633.54	2.33
	2396	vmiprvse	194.73	125.37	1.33	69.36	1.33
Jan 03, 2008 05:24:01							
	1032	svchost	770.82	137.36	3	633.46	2.33
	2396	vmiprvse	194.7	125.36	1.33	69.35	1.33
	628	csrss	16	16	0.67	0	0
Jan 03, 2008 05:14:10							
	4	System	45069.31	45069.31	88.03	0	0
	1032	svchost	20346.68	7702.94	9.67	12643.73	1.33
	828	cmd	73.69	73.69	1.67	0	0
	2876	tomcat	16.34	0	0	16.34	0.67
	628	csrss	16	16	1.33	0	0

Figure 2.14: The detailed diagnosis of the Percent virtual busy measure

2.6.1.1 Configuring Users for VM Monitoring

In order to enable the eG agent to connect to VMs in multiple domains and pull out metrics from them, the eG administrative interface provides a special page using which the different **DOMAIN** names, and their corresponding **ADMIN USER** names and **ADMIN PASSWORDS** can be specified. To access this page, just click on the **Click here** hyperlink in any of the VM test configuration pages.

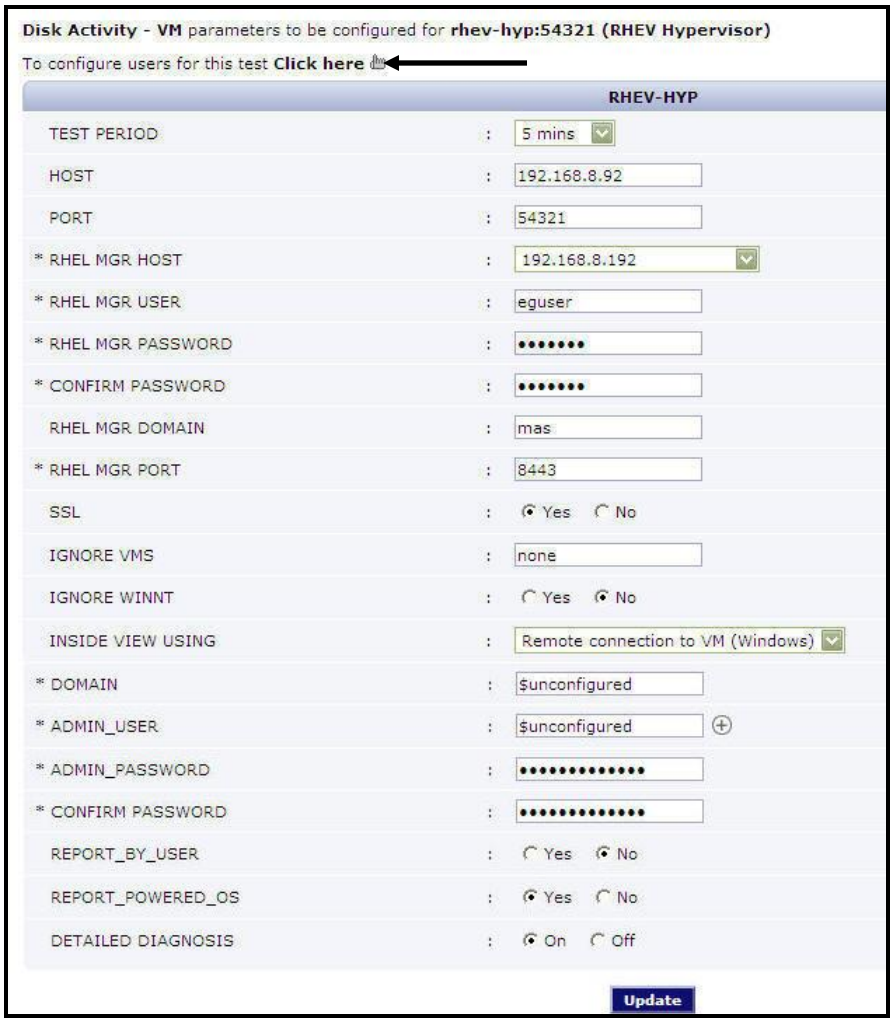

Figure 2.15: Configuring a VM test

Upon clicking, Figure 2.16 will appear, using which the VM user details can be configured.

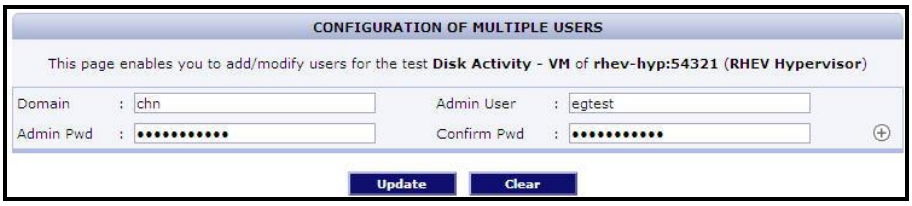

Figure 2.16: The VM user configuration page

To add a user specification, do the following:

1. First, provide the name of the **Domain** to which the VMs belong (see Figure 2.16). If one/more VMs do not belong to any domain, then, specify *none* here.
2. The eG agent must be configured with user privileges that will allow the agent to communicate with the VMs in a particular domain and extract statistics. If *none* is specified against **Domain**, then a local user account can be provided against **Admin User**. On the other hand, if a valid **Domain** name has been specified, then a domain administrator account can be provided in the **Admin User** text box. If key-based authentication is implemented between the eG agent and the SSH daemon of a

Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

3. The password of the specified **Admin User** should be mentioned in the **Admin Pwd** text box.
4. Confirm the password by retyping it in the **Confirm Pwd** text box.
5. To add more users, click on the ⊕ button in Figure 2.16. This will allow you to add one more user specification as depicted by Figure 2.17.

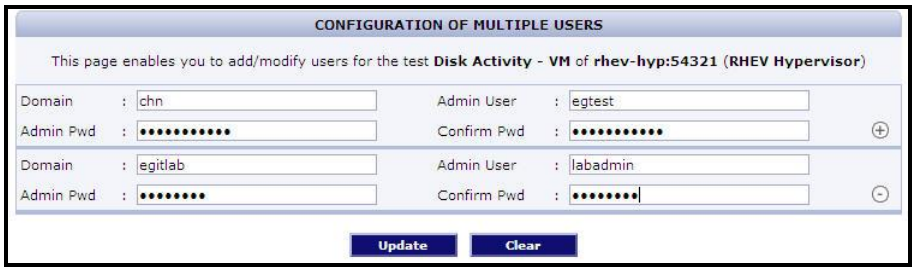

Figure 2.17: Adding another user

6. In some virtualized environments, the same **Domain** could be accessed using multiple **Admin User** names. For instance, to login to a **Domain** named *egitlab*, the eG agent can use the **Admin User** name *labadmin* or the **Admin User** name *jadmn*. You can configure the eG agent with the credentials of both these users as shown by Figure 2.18.

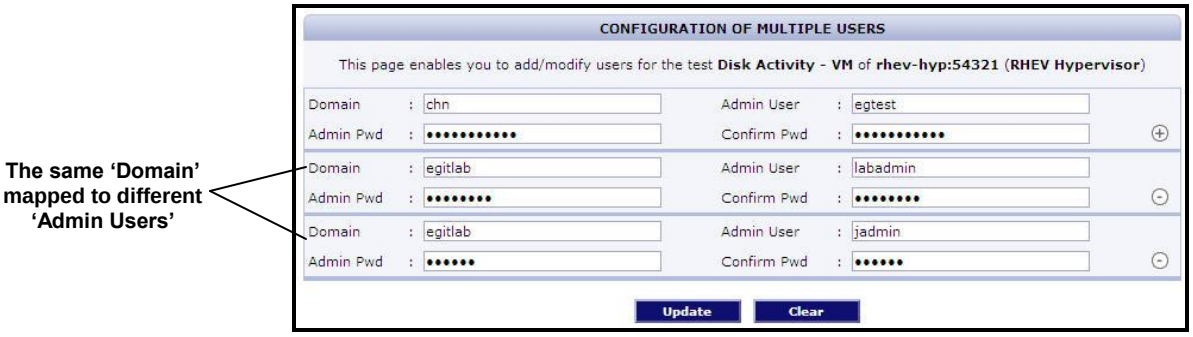

Figure 2.18: Associating a single domain with different admin users

When this is done, then, while attempting to connect to the domain, the eG agent will begin by using the first **Admin User** name of the specification. In the case of Figure 2.18, this will be *labadmin*. If, for some reason, the agent is unable to login using the first **Admin User** name, then it will try to login again, but this time using the second **Admin User** name of the specification - i.e., *jadmn* in our example (see Figure 2.18). If the first login attempt itself is successful, then the agent will ignore the second **Admin User** name.

7. To clear all the user specifications, simply click the **Clear** button in Figure 2.18.
8. To remove the details of a particular user alone, just click the ⊖ button in Figure 2.18.

9. To save the specification, just click on the **Update** button in Figure 2.18. This will lead you back to the test configuration page, where you will find the multiple domain names, user names, and passwords listed against the respective fields (see Figure 2.18).

Figure 2.19: The test configuration page displaying multiple domain names, user names, and passwords

2.6.2 Disk Space - VM Test

This test monitors the space usage of every disk partition on a guest.

Purpose	To measure the space usage of every disk partition on each guest of an RHEV server
Target of the test	An RHEV hypervisor
Agent deploying the test	A remote 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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

8. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

9. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

   Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights.** Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

10. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

    - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
    - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

      Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

      If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

<table>
<tr><td></td><td colspan="3">

If the guests belong to different domains - In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple DOMAIN names, multiple ADMIN USER names and ADMIN PASSWORDS would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page.

To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document.
If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' - In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to none.

11. REPORT BY USER - While monitoring a RHEV Hypervisor, the REPORT BY USER flag is set to NO by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a RHEV Hypervisor - VDI, this flag is set to YES by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every username_on_virtualmachinename.
12. REPORT POWERED OS - This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.

If the REPORT POWERED OS flag is set to Yes (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their virtual machine name and not by the username_on_virtualmachinename. On the other hand, if the REPORT POWERED OS flag is set to No, then this test will not report measures for those VMs to which no users are logged in currently.
</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every combination of virtual_guest:disk_partition</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Total capacity:
Indicates the total capacity of a disk partition.</td><td>MB</td><td></td></tr>
<tr><td></td><td>Used space:
Indicates the amount of space used in a disk partition.</td><td>MB</td><td></td></tr>
</table>

	Free space: Indicates the current free space available for each disk partition of a system.	MB	
	Percent usage: Indicates the percentage of space usage on each disk partition of a system.	Percent	A value close to 100% can indicate a potential problem situation where applications executing on the guest may not be able to write data to the disk partition(s) with very high usage.

2.6.3 Memory Usage - VM Test

This test reports statistics related to the usage of physical memory of the VMs.

Purpose	Reports statistics related to the usage of physical memory of the VMs
Target of the test	An RHEV Hypervisor
Agent deploying the test	A remote 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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

8. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

9. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

   Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights.** Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

10. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

    - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
    - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

      Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

      If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

<table>
<tr><td></td><td colspan="3">

If the guests belong to different domains - In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple DOMAIN names, multiple ADMIN USER names and ADMIN PASSWORDS would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page.

To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document.
If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' - In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to none.

11. REPORT BY USER - While monitoring a RHEV Hypervisor, the REPORT BY USER flag is set to NO by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a RHEV Hypervisor - VDI, this flag is set to YES by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every username_on_virtualmachinename.
12. REPORT POWERED OS - This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.

If the REPORT POWERED OS flag is set to Yes (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their virtual machine name and not by the username_on_virtualmachinename. On the other hand, if the REPORT POWERED OS flag is set to No, then this test will not report measures for those VMs to which no users are logged in currently.
13. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG 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.

</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every VM on an RHEV server</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

	Total physical memory: Indicates the total physical memory of this VM.	MB	
	Used physical memory: Indicates the used physical memory of this VM.	MB	
	Free physical memory: Indicates the free physical memory of the VM.	MB	This measure typically indicates the amount of memory available for use by applications running on the target VM. On Unix operating systems (AIX and Linux), the operating system tends to use parts of the available memory for caching files, objects, etc. When applications require additional memory, this is released from the operating system cache. Hence, to understand the true free memory that is available to applications, the eG agent reports the sum of the free physical memory and the operating system cache memory size as the value of the *Free physical memory* measure while monitoring AIX and Linux guest operating systems.
	Physical memory utilized: Indicates the percent usage of physical memory by this VM.	Percent	Ideally, the value of this measure should be low. While sporadic spikes in memory usage could be caused by one/more rogue processes on the VM, a consistent increase in this value could be a cause for some serious concern, as it indicates a gradual, but steady erosion of valuable memory resources. If this unhealthy trend is not repaired soon, it could severely hamper VM performance, causing anything from a slowdown to a complete system meltdown. You can use the detailed diagnosis of this measure to figure out which processes on the VM are consuming memory excessively.

	Available physical memory: Indicates the amount of physical memory, immediately available for allocation to a process or for system use.	MB	Not all of the *Available physical memory* is *Free physical memory*. Typically, *Available physical memory* is made up of the Standby List, Free List, and Zeroed List. When Windows wants to trim a process' working set, the trimmed pages are moved (usually) to the Standby List. From here, they can be brought back to life in the working set with only a soft page fault (much faster than a hard fault, which would have to talk to the disk). If a page stays in the standby List for a long time, it gets freed and moved to the Free List. In the background, there is a low priority thread (actually, the only thread with priority 0) which takes pages from the Free List and zeros them out. Because of this, there is usually very little in the Free List. All new allocations always come from the Zeroed List, which is memory pages that have been overwritten with zeros. This is a standard part of the OS' cross-process security, to prevent any process ever seeing data from another. If the Zeroed List is empty, Free List memory is zeroed and used or, if that is empty too, Standby List memory is freed, zeroed, and used. It is because all three can be used with so little effort that they are all counted as "available". A high value is typically desired for this measure. **This measure will be available for Windows 2008 VMs only.**

	Modified memory: Indicates the amount of memory that is allocated to the modified page list.	MB	This memory contains cached data and code that is not actively in use by processes, the system and the system cache. This memory needs to be written out before it will be available for allocation to a process or for system use. Cache pages on the modified list have been altered in memory. No process has specifically asked for this data to be in memory, it is merely there as a consequence of caching. Therefore it can be written to disk at any time (not to the page file, but to its original file location) and reused. However, since this involves I/O, it is not considered to be Available physical memory. **This measure will be available for Windows 2008 VMs only.**
	Standby memory: Indicates the amount of memory assigned to the standby list.	MB	This memory contains cached data and code that is not actively in use by processes, the system and the system cache. It is immediately available for allocation to a process or for system use. If the system runs out of available free and zero memory, memory on lower priority standby cache page lists will be repurposed before memory on higher priority standby cache page lists. Typically, Standby memory is the aggregate of Standby Cache Core Bytes,Standby Cache Normal Priority Bytes, and Standby Cache Reserve Bytes. Standby Cache Core Bytes is the amount of physical memory, that is assigned to the core standby cache page lists. Standby Cache Normal Priority Bytes is the amount of physical memory, that is assigned to the normal priority standby cache page lists. Standby Cache Reserve Bytes is the amount of physical memory, that is assigned to the reserve standby cache page lists. **This measure will be available for Windows 2008 VMs only.**

	Cached memory: This measure is an aggregate of *Standby memory* and *Modified memory*.	MB	**This measure will be available for Windows 2008 VMs only.**

Note:

While monitoring Linux/AIX guest operating systems, you may observe discrepancies between the value of the *Physical memory utilized* measure and the memory usage percentages reported per process by the detailed diagnosis of the same measure. This is because, while the *Physical memory utilized* measure takes into account the memory in the OS cache of the Linux/AIX VM, the memory usage percent that the detailed diagnosis reports per process does not consider the OS cache memory.

2.6.4 System Details - VM Test

This test collects various metrics pertaining to the CPU and memory usage of every processor supported by a guest. The details of this test are as follows:

Purpose	To measure the CPU and memory usage of each guest of an RHEV Hypervisor
Target of the test	An RHEV Hypervisor
Agent deploying the test	A remote 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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

8. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

9. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

   Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights.** Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

10. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

    - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
    - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

      Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

      If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

- **If the guests belong to different domains -** In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple **DOMAIN** names, multiple **ADMIN USER** names and **ADMIN PASSWORDS** would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page.

  To access this page, simply click on the **Click here** hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document.

- **If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' -** In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to *none*.

11. **REPORT BY USER** - While monitoring a *RHEV Hypervisor*, the **REPORT BY USER** flag is set to **NO** by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a *RHEV Hypervisor - VDI*, this flag is set to **YES** by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every *username_on_virtualmachinename*.

12. **REPORT POWERED OS** - **This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.**

    If the **REPORT POWERED OS** flag is set to **Yes** (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their *virtual machine name* and not by the *username_on_virtualmachinename*. On the other hand, if the **REPORT POWERED OS** flag is set to **No**, then this test will not report measures for those VMs to which no users are logged in currently.

13. **USE TOP FOR DD** – **This flag is only applicable to Linux VMs.** By default, this parameter is set to **No**. This indicates that, by default, this test will report the detailed diagnosis of the *Virtual CPU utilization* measure for each processor on a Linux VM by executing the *usr/bin/ps* command. On some Linux flavors however, this command may not function properly. In such cases, set the **USE TOP FOR DD** parameter to **Yes**. This will enable the eG agent to extract the detailed diagnosis of the *Virtual CPU utilization* measure by executing the */usr/bin/top* command instead.

14. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG 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 every combination of *virtual_guest:processor*		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Virtual CPU utilization: This measurement indicates the percentage of CPU utilized by the processor.	Percent	A high value could signify a CPU bottleneck. The CPU utilization may be high because a few processes are consuming a lot of CPU, or because there are too many processes contending for a limited resource. The detailed diagnosis of this test reveals the top-10 CPU-intensive processes on the guest.
	System usage of virtual CPU: Indicates the percentage of CPU time spent for system-level processing.	Percent	An unusually high value indicates a problem and may be due to too many system-level tasks executing simultaneously.
	Run queue in VM: Indicates the instantaneous length of the queue in which threads are waiting for the processor cycle. This length does not include the threads that are currently being executed.	Number	A value consistently greater than 2 indicates that many processes could be simultaneously contending for the processor.
	Blocked processes in VM: Indicates the number of processes blocked for I/O, paging, etc.	Number	A high value could indicate an I/O problem on the guest (e.g., a slow disk).
	Swap memory in VM: Denotes the committed amount of virtual memory. This corresponds to the space reserved for virtual memory on disk paging file(s).	MB	An unusually high value for the swap usage can indicate a memory bottleneck. Check the memory utilization of individual processes to figure out the process(es) that has (have) maximum memory consumption and look to tune their memory usages and allocations accordingly.

	Free memory in VM: Indicates the free memory available.	MB	This measure typically indicates the amount of memory available for use by applications running on the target VM. On Unix operating systems (AIX and Linux), the operating system tends to use parts of the available memory for caching files, objects, etc. When applications require additional memory, this is released from the operating system cache. Hence, to understand the true free memory that is available to applications, the eG agent reports the sum of the free physical memory and the operating system cache memory size as the value of the *Free memory in VM* measure while monitoring AIX and Linux guest operating systems. The detailed diagnosis of this measure, if enabled, lists the top 10 processes responsible for maximum memory consumption on the target VM.
	Scan rate in VM: Indicates the memory scan rate.	Pages/Sec	A high value is indicative of memory thrashing. Excessive thrashing can be detrimental to guest performance.

Note:

For multi-processor systems, where the CPU statistics are reported for each processor on the system, the statistics that are system-specific (e.g., run queue length, free memory, etc.) are only reported for the "Summary" descriptor of this test.

The detailed diagnosis capability of the *Virtual CPU utilization* measure, if enabled, provides a listing of the top 10 CPU-consuming processes (see Figure 2.20). In the event of a Cpu bottleneck, this information will enable users to identify the processes consuming a high percentage of CPU time. The users may then decide to stop such processes, so as to release the CPU resource for more important processing purposes.

Lists the top 10 CPU processes

Time	PID	%CPU	ARGS
Jan 03, 2008 05:43:18	4	0.52	system
Jan 03, 2008 05:32:51	4	0.52	system
Jan 03, 2008 05:22:20	4	0.52	system
Jan 03, 2008 05:12:21	4	0.52	system
Jan 03, 2008 05:02:47	4	0.52	system
Jan 03, 2008 05:02:47	1768	0.52	xenservice
Jan 03, 2008 04:53:13	4	0.52	system

Figure 2.20: The top 10 CPU consuming processes

Note:

While instantaneous spikes in CPU utilization are captured by the eG agents and displayed in the Measures page, the detailed diagnosis will not capture/display such instantaneous spikes. Instead, detailed diagnosis will display only a consistent increase in CPU utilization observed over a period of time.

The detailed diagnosis of the *Free memory in VM* measure, if enabled, lists the top 10 processes responsible for maximum memory consumption on the guest (see Figure 2.21). This information will enable administrators to identify the processes that are causing the depletion in the amount of free memory on the host. The administrators can then decide to kill such expensive processes.

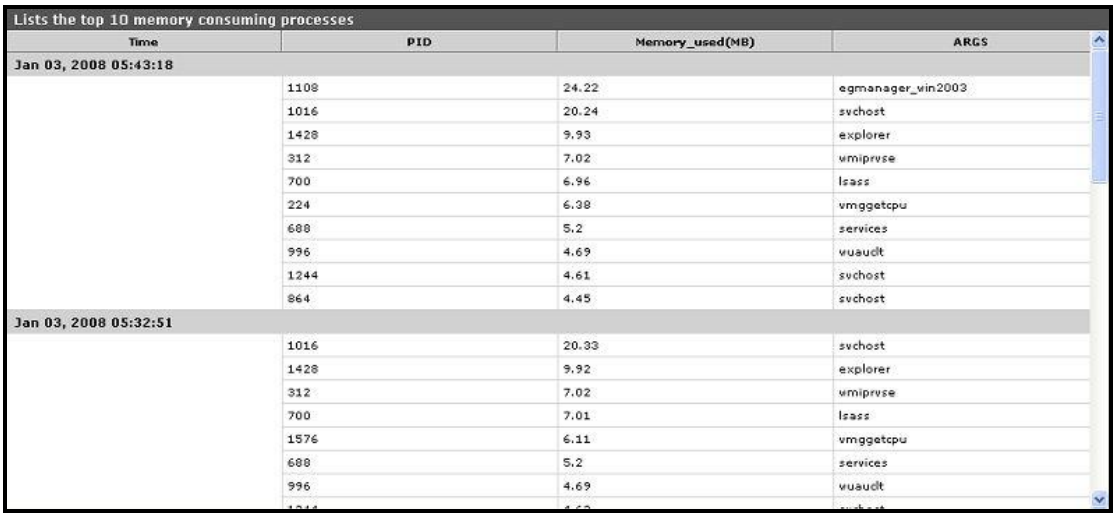

Lists the top 10 memory consuming processes

Time	PID	Memory_used(MB)	ARGS
Jan 03, 2008 05:43:18			
	1108	24.22	egmanager_win2003
	1016	20.24	svchost
	1428	9.93	explorer
	312	7.02	vmiprvse
	700	6.96	lsass
	224	6.38	vmggetcpu
	688	5.2	services
	996	4.69	vuaudit
	1244	4.61	svchost
	864	4.45	svchost
Jan 03, 2008 05:32:51			
	1016	20.33	svchost
	1428	9.92	explorer
	312	7.02	vmiprvse
	700	7.01	lsass
	1576	6.11	vmggetcpu
	688	5.2	services
	996	4.69	vuaudit

Figure 2.21: The detailed diagnosis of the Free memory measure listing the top 10 memory consuming processes

2.6.5 Uptime - VM Test

In most virtualized environments, it is essential to monitor the uptime of VMs hosting critical server applications in the infrastructure. By tracking the uptime of each of the VMs, administrators can determine what percentage of time a VM has been up. Comparing this value with service level targets, administrators can determine the most trouble-prone areas of the virtualized infrastructure.

In some environments, administrators may schedule periodic reboots of their VM. By knowing that a specific VM has been up for an unusually long time, an administrator may come to know that the scheduled reboot task is not working on a VM.

The Uptime - Guest test included in the eG agent monitors the uptime of each VM on an RHEV Hypervisor.

Purpose	To monitor the uptime of each guest on a
Target of the test	An RHEV Hypervisor
Agent deploying the test	A remote 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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

8. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

9. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights**. Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

10. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

    - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
    - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

      Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

      If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

	➢ **If the guests belong to different domains -** In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple **DOMAIN** names, multiple **ADMIN USER** names and **ADMIN PASSWORDS** would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the **Click here** hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document. ➢ **If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' -** In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to *none*. 11. **REPORT BY USER** - While monitoring a *RHEV Hypervisor*, the **REPORT BY USER** flag is set to **NO** by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a *RHEV Hypervisor - VDI*, this flag is set to **YES** by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every *username_on_virtualmachinename*. 12. **REPORT POWERED OS** - **This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.** If the **REPORT POWERED OS** flag is set to **Yes** (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their *virtual machine name* and not by the *username_on_virtualmachinename*. On the other hand, if the **REPORT POWERED OS** flag is set to **No**, then this test will not report measures for those VMs to which no users are logged in currently. 13. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG 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 every guest on an RHEV Hypervisor		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Has VM been rebooted**: Indicates whether the VM has been rebooted during the last measurement period or not.	Boolean	If this measure shows 1, it means that the guest was rebooted during the last measurement period. By checking the time periods when this metric changes from 0 to 1, an administrator can determine the times when this guest was rebooted.
	Uptime of VM during the last measure period: Indicates the time period that the VM has been up since the last time this test ran.	Secs	If the guest has not been rebooted during the last measurement period and the agent has been running continuously, this value will be equal to the measurement period. If the guest was rebooted during the last measurement period, this value will be less than the measurement period of the test. For example, if the measurement period is 300 secs, and if the guest was rebooted 120 secs back, this metric will report a value of 120 seconds. The accuracy of this metric is dependent on the measurement period - the smaller the measurement period, greater the accuracy.
	Total uptime of the VM: Indicates the total time that the VM has been up since its last reboot.	Mins	Administrators may wish to be alerted if a guest has been running without a reboot for a very long period. Setting a threshold for this metric allows administrators to determine such conditions.

Note:

If a value less than a minute is configured as the **TEST PERIOD** of the Uptime - Guest test, then, the **Uptime during the last measure period** measure will report the value 0 for Linux VMs (only) until the minute boundary is crossed. For instance, if you configure the Uptime - Guest test to run every 10 seconds, then, for the first 5 test execution cyles (i.e., 10 x 5 = 50 seconds), the **Uptime during the last measure period** measure will report the value 0 for Linux VMs; however, the sixth time the test executes (i.e, when test execution touches the 1 minute boundary), this measure will report the value 60 seconds for the same VMs. Thereafter, every sixth measurement period will report 60 seconds as the uptime of the Linux VMs. This is because, Linux operating systems report uptime only in minutes and not in seconds.

2.6.6 Windows Memory - VM Test

To understand the metrics reported by this test, it is essential to understand how memory is handled by the operating system. On any Windows system, memory is partitioned into a part that is available

for user processes, and another that is available to the OS kernel. The kernel memory area is divided into several parts, with the two major parts (called "pools") being a nonpaged pool and a paged pool. The nonpaged pool is a section of memory that cannot, under any circumstances, be paged to disk. The paged pool is a section of memory that can be paged to disk. (Just being stored in the paged pool doesn't necessarily mean that something has been paged to disk. It just means that it has either been paged to disk or it could be paged to disk.) Sandwiched directly in between the nonpaged and paged pools (although technically part of the nonpaged pool) is a section of memory called the "System Page Table Entries," or "System PTEs." The WindowsMemory - Guest test tracks critical metrics corresponding to the System PTEs and the pool areas of kernel memory of a Windows virtual machine.

Purpose	Tracks critical metrics corresponding to the System PTEs and the pool areas of kernel memory of each Windows-based virtual guest of an RHEV Hypervisor
Target of the test	An RHEV Hypervisor
Agent deploying the test	A remote 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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

8. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

9. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

   Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights**. Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

10. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

    - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
    - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

      Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

      If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

<table>
<tr><td></td><td colspan="3">

If the guests belong to different domains - In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple DOMAIN names, multiple ADMIN USER names and ADMIN PASSWORDS would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page.

To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document.
If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' - In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to none.

11. REPORT BY USER - While monitoring a RHEV Hypervisor, the REPORT BY USER flag is set to NO by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a RHEV Hypervisor - VDI, this flag is set to YES by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every username_on_virtualmachinename.
12. REPORT POWERED OS - This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.

If the REPORT POWERED OS flag is set to Yes (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their virtual machine name and not by the username_on_virtualmachinename. On the other hand, if the REPORT POWERED OS flag is set to No, then this test will not report measures for those VMs to which no users are logged in currently.
</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every Windows VM guest/user on the monitored RHEV hypervisor</td></tr>
<tr><td>Measurements made by the</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

test	**Free entries in system page table:** Indicates the number of page table entries not currently in use by the guest.	Number	The maximum number of System PTEs that a server can have is set when the server boots. In heavily-used servers, you can run out of system PTEs. You can use the registry to increase the number of system PTEs, but that encroaches into the paged pool area, and you could run out of paged pool memory. Running out of either one is bad, and the goal should be to tune your server so that you run out of both at the exact same time. Typically, the value of this metric should be above 3000.
	Page read rate in VM: Indicates the average number of times per second the disk was read to resolve hard fault paging.	Reads/Sec	
	Page write rate in VM: Indicates the average number of times per second the pages are written to disk to free up the physical memory.	Writes/Sec	
	Page input rate in VM: Indicates the number of times per second that a process needed to access a piece of memory that was not in its working set, meaning that the guest had to retrieve it from the page file.	Pages/Sec	

	Page output rate in VM: Indicates the number of times per second the guest decided to trim a process's working set by writing some memory to disk in order to free up physical memory for another process.	Pages/Sec	This value is a critical measure of the memory utilization on a guest. If this value never increases, then there is sufficient memory in the guest. Instantaneous spikes of this value are acceptable, but if the value itself starts to rise over time or with load, it implies that there is a memory shortage on the guest.
	Memory pool non-paged data in VM: Indicates the total size of the kernel memory nonpaged pool.	MB	The kernel memory nonpage pool is an area of guest memory (that is, memory used by the guest operating system) for kernel objects that cannot be written to disk, but must remain in memory as long as the objects are allocated. Typically, there should be no more than 100 MB of non-paged pool memory being used.
	Memory pool paged data in VM : Indicates the total size of the Paged Pool.	MB	If the Paged Pool starts to run out of space (when it's 80% full by default), the guest will automatically take some memory away from the System File Cache and give it to the Paged Pool. This makes the System File Cache smaller. However, the system file cache is critical, and so it will never reach zero. Hence, a significant increase in the paged pool size is a problem. This metric is a useful indicator of memory leaks in a guest. A memory leak occurs when the guest allocates more memory to a process than the process gives back to the pool. Any time of process can cause a memory leak. If the amount of paged pool data keeps increasing even though the workload on the guest remains constant, it is an indicator of a memory leak.

2.6.7 Windows Network Traffic - VM Test

This is an internal test that monitors the incoming and outgoing traffic through each Windows guest of an RHEV Hypervisor.

Purpose	To measure the incoming and outgoing traffic through each Windows-based guest of an RHEV Hypervisor
Target of the test	An RHEV Hypervisor

Agent deploying the test	A remote 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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

8. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

9. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

   Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights**. Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

10. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

    - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
    - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

      Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

      If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

<table>
<tr><td></td><td colspan="3">

If the guests belong to different domains - In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple DOMAIN names, multiple ADMIN USER names and ADMIN PASSWORDS would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page.

To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document.
If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' - In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to none.

11. REPORT BY USER - While monitoring a RHEV Hypervisor, the REPORT BY USER flag is set to NO by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a RHEV Hypervisor - VDI, this flag is set to YES by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every username_on_virtualmachinename.
12. REPORT POWERED OS - This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.

If the REPORT POWERED OS flag is set to Yes (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their virtual machine name and not by the username_on_virtualmachinename. On the other hand, if the REPORT POWERED OS flag is set to No, then this test will not report measures for those VMs to which no users are logged in currently.
</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every Windows_virtual_guest:network_interface combination combination</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Incoming traffic:

Indicates the rate at which data (including framing characters) is received on a network interface.</td><td>Mbps</td><td>An abnormally high rate of incoming traffic may require additional analysis.</td></tr>
<tr><td></td><td>Outgoing traffic: Represents the rate at which data (including framing characters) is sent on a network interface.</td><td>Mbps</td><td>An abnormally high rate of outgoing traffic may require additional analysis.</td></tr>
</table>

	Maximum bandwidth: An estimate of the capacity of a network interface.	Mbps	
	Bandwidth usage: Indicates the percentage of bandwidth used by a network interface.	Percent	By comparing the bandwidth usage with the maximum bandwidth of an interface, an administrator can determine times when the network interface is overloaded or is being a performance bottleneck.
	Output queue length: Indicates the length of the output packet queue (in packets)	Number	If this is longer than 2, delays are being experienced and the bottleneck should be found and eliminated if possible.
	Outbound packet errors: The number of outbound packets that could not be transmitted because of errors	Number	Ideally, number of outbound errors should be 0.
	Inbound packet errors: The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.	Number	Ideally, number of inbound errors should be 0.

If the WindowsNetTraffic - Guest test is not reporting measures for a guest, make sure that you have enabled the SNMP service for the guest.

2.6.8 Network Traffic - VM Test

This is an internal test that monitors the incoming and outgoing traffic through each Linux guest on an RHEV Hypervisor.

Purpose	To measure the incoming and outgoing traffic through each Linux guest on an RHEV Hypervisor
Target of the test	An RHEV Hypervisor
Agent deploying the test	A remote 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. **RHEL MGR HOST**, **RHEL MGR PORT**, **RHEL MGR DOMAIN**, **RHEL MGR USER**, **RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

8. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

9. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

   Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights**. Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

10. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

    - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
    - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

      Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

      If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

	➢ **If the guests belong to different domains -** In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple **DOMAIN** names, multiple **ADMIN USER** names and **ADMIN PASSWORDS** would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the **Click here** hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document. ➢ **If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' -** In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to *none*. 11. **REPORT BY USER** - While monitoring a *RHEV Hypervisor*, the **REPORT BY USER** flag is set to **NO** by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a *RHEV Hypervisor - VDI*, this flag is set to **YES** by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every *username_on_virtualmachinename*. 12. **REPORT POWERED OS** - **This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.** If the **REPORT POWERED OS** flag is set to **Yes** (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their *virtual machine name* and not by the *username_on_virtualmachinename*. On the other hand, if the **REPORT POWERED OS** flag is set to **No**, then this test will not report measures for those VMs to which no users are logged in currently.
Outputs of the test	One set of results for every *Linux virtual_guest:network_interface* combination

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Incoming traffic: Indicates the rate of incoming traffic.	Pkts/Sec	An increase in traffic to the guest can indicate an increase in accesses to the guest (from users or from other applications) or that the guest is under an attack of some form.
	Outgoing traffic: Represents the rate of outgoing traffic.	Pkts/Sec	An increase in traffic from the guest can indicate an increase in accesses to the guest (from users or from other applications).

2.6.9 Tcp - VM Test

This test tracks various statistics pertaining to TCP connections to and from each guest of an RHEV Hypervisor.

Purpose	To measure statistics pertaining to the TCP layer of a guest
Target of the test	An RHEV Hypervisor
Agent deploying the test	A remote 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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

8. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

9. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

   Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights**. Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

10. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

    - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
    - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

      Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

      If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

<table>
<tr><td></td><td>

If the guests belong to different domains - In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple DOMAIN names, multiple ADMIN USER names and ADMIN PASSWORDS would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page.

To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document.
If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' - In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to none.

11. REPORT BY USER - While monitoring a RHEV Hypervisor, the REPORT BY USER flag is set to NO by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a RHEV Hypervisor - VDI, this flag is set to YES by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every username_on_virtualmachinename.
12. REPORT POWERED OS - This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.

If the REPORT POWERED OS flag is set to Yes (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their virtual machine name and not by the username_on_virtualmachinename. On the other hand, if the REPORT POWERED OS flag is set to No, then this test will not report measures for those VMs to which no users are logged in currently.
</td></tr>
<tr><td>Outputs of the test</td><td>One set of results for each powered-on guest on the RHEV Hypervisor being monitored</td></tr>
<tr><td>Measurements made by the test</td><td>
<table>
<tr><th>Measurement</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Incoming connections to VM:
Indicates the connections per second received by the guest.</td><td>Conns/Sec</td><td>A high value can indicate an increase in input load.</td></tr>
<tr><td>Outgoing connections to VM:
Indicates the connections per second initiated by the guest.</td><td>Conns/Sec</td><td>A high value can indicate that one or more of the applications executing on the guest have started using a number of TCP connections to some other guest or host.</td></tr>
</table>
</td></tr>
</table>

	Current connections to VM: Indicates the currently established connections.	Number	A sudden increase in the number of connections established on a guest can indicate either an increase in load to one or more of the applications executing on the guest, or that one or more of the applications are experiencing a problem (e.g., a slow down). On Microsoft Windows, the current connections metrics is the total number of TCP connections that are currently in the ESTABLISHED or CLOSE_WAIT states.
	Connection drops on VM: Indicates the rate of established TCP connections dropped from the TCP listen queue.	Conns/Sec	This value should be 0 for most of the time. Any non-zero value implies that one or more applications on the guest are under overload.
	Connection failures on VM: Indicates the rate of half open TCP connections dropped from the listen queue.	Conns/Sec	This value should be 0 for most of the time. A prolonged non-zero value can indicate either that the server is under SYN attack or that there is a problem with the network link to the server that is resulting in connections being dropped without completion.

2.6.10 Tcp Traffic - VM Test

Since most popular applications rely on the TCP protocol for their proper functioning, traffic monitoring at the TCP protocol layer can provide good indicators of the performance seen by the applications that use TCP. The most critical metric at the TCP protocol layer is the percentage of retransmissions. Since TCP uses an exponential back-off algorithm for its retransmissions, any retransmission of packets over the network (due to network congestion, noise, data link errors, etc.) can have a significant impact on the throughput seen by applications that use TCP. This test monitors the TCP protocol traffic to and from a guest, and particularly monitors retransmissions.

Purpose	Monitors the TCP protocol traffic to and from each guest of an RHEV Hypervisor, and particularly measures the percentage of retransmission
Target of the test	An RHEV Hypervisor
Agent deploying the test	A remote 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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

8. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

9. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

   Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights**. Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

10. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

    - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
    - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

      Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

      If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

	➢ **If the guests belong to different domains -** In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple **DOMAIN** names, multiple **ADMIN USER** names and **ADMIN PASSWORDS** would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the **Click here** hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document. ➢ **If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' -** In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to *none*. 11. **REPORT BY USER** - While monitoring a *RHEV Hypervisor*, the **REPORT BY USER** flag is set to **NO** by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a *RHEV Hypervisor - VDI*, this flag is set to **YES** by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every *username_on_virtualmachinename*. 12. **REPORT POWERED OS** - **This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.** If the **REPORT POWERED OS** flag is set to **Yes** (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their *virtual machine name* and not by the *username_on_virtualmachinename*. On the other hand, if the **REPORT POWERED OS** flag is set to **No**, then this test will not report measures for those VMs to which no users are logged in currently. 13. **SEGMENTS SENT MIN** - Specify the minimum threshold for the number of segments sent/transmitted over the network. The default value is *none*; in this case, the test will not compute/report the **Retransmit ratio from VM** measure. On the other hand, if a non-zero value is provided against the **SEGMENTS SENT MIN** parameter, then the test will begin computing the **Retransmit ratio from VM** measure only when the value of the **Segments sent by VM** measure crosses the value set for this parameter. This is done to ensure that no false alerts are generated by the eG Enterprise system for the **Retransmit ratio from VM** measure.		
Outputs of the test	One set of results for each powered-on guest on the RHEV Hypervisor being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Segments received by VM:** Indicates the rate at which segments are received by the guest.	Segments/Sec	
	Segments sent by VM: Indicates the rate at which segments are sent to clients or other guests	Segments/Sec	
	Retransmits by VM: Indicates the rate at which segments are being retransmitted by the guest	Segments/Sec	
	Retransmit ratio from VM: Indicates the ratio of the rate of data retransmissions to the rate of data being sent by the guest	Percent	Ideally, the retransmission ratio should be low (< 5%). Most often retransmissions at the TCP layer have significant impact on application performance. Very often a large number of retransmissions are caused by a congested network link, bottlenecks at a router causing buffer/queue overflows, or by lousy network links due to poor physical layer characteristics (e.g., low signal to noise ratio). By tracking the percentage of retransmissions at a guest, an administrator can quickly be alerted to problem situations in the network link(s) to the guest that may be impacting the service performance.

2.6.11 Handles Usage - VM Test

This test monitors and tracks the handles opened by processes running in a target Windows virtual machine.

Purpose	Monitors and tracks the handles opened by processes running in a target Windows virtual machine
Target of the test	An RHEV Hypervisor
Agent deploying the test	A remote 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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN, RHEL MGR USER, RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

8. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

9. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

   Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights**. Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

10. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

    - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
    - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

      Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

      If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

<table>
<tr><td></td><td colspan="3">

If the guests belong to different domains - In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple DOMAIN names, multiple ADMIN USER names and ADMIN PASSWORDS would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page.

To access this page, simply click on the Click here hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document.
If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' - In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to none.

11. REPORT BY USER - While monitoring a RHEV Hypervisor, the REPORT BY USER flag is set to NO by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a RHEV Hypervisor - VDI, this flag is set to YES by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every username_on_virtualmachinename.
12. REPORT POWERED OS - This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.

If the REPORT POWERED OS flag is set to Yes (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their virtual machine name and not by the username_on_virtualmachinename. On the other hand, if the REPORT POWERED OS flag is set to No, then this test will not report measures for those VMs to which no users are logged in currently.
</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each powered-on guest on the RHEV Hypervisor monitored</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Handles used by processes:

Indicates the number of handles opened by various processes running in a target Windows virtual machine in the last measurement period.</td><td>Number</td><td>Use the detailed diagnosis of this measure to determine the top-10 processes in terms of number of handles opened. This information brings to light those processes with too many open handles. By closely tracking the handle usage of these processes over time, you can identify potential handle leaks.</td></tr>
</table>

	Processes using handles above limit in the VM: Indicates the number of processes that have opened the handles on or above the value defined in the input parameter - **HANDLES GROWTH LIMIT**.	Number	Using the detailed diagnosis of this measure, you can accurately isolate the process(es) that has opened more handles than the permitted limit. A high value of this measure indicates that too many processes are opening handles excessively. You might want to closely observe the handle usage of these processes over time to figure out whether the spike in usage is sporadic or consistent. A consistent increase in handle usage could indicate a handle leak.

The detailed diagnosis of the *Handles used by processes* measure, if enabled, lists the names of top-10 processes in terms of handle usage, the number of handles each process uses, the process ID, and the ID of the parent process.

List of top 10 processes in a VM that are holding handles

Time	Process Name	Handles used	Process ID	Parent PID
Jan 29, 2009 12:00:49				
	System	3359	0	4
	js	1718	540	6420
	svchost	1208	540	1012
	lsass	1112	492	552
	csrss	1097	420	468
	winlogon	564	420	492
	ImaSrv	559	540	3696
	Rtvscan	536	540	3936
	tomcat	485	540	6572
	services	482	492	540

Figure 2.22: The detailed diagnosis of the Handles used by processes measure

The detailed diagnosis of the *Processes using handles above limit in VM* measure, if enabled, lists the details of processes that are using more handles than the configured limit.

List of processes in a VM that are using handles above the configured handle growth value

Time	Process Name	Handles used	Process ID	Parent PID
Jan 29, 2009 17:54:18	eGRSvc	62410	412	11512

Figure 2.23: The detailed diagnosis of the Processes using handles above limit in VM measure

2.6.12 Windows Services - VM Test

This test tracks the status (whether running or have stopped) of services executing on Windows virtual machines.

Purpose	Tracks the status (whether running or have stopped) of services executing on Windows virtual machines
Target of the test	An RHEV Hypervisor
Agent deploying the	A remote agent

test	
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. **RHEL MGR HOST, RHEL MGR PORT, RHEL MGR DOMAIN**, **RHEL MGR USER**, **RHEL MGR PASSWORD** - To auto-discover the VMs on a target RHEV hypervisor and obtain the *outside view* of the performance of each VM, the eG agent needs to connect to the RHEV Manager that manages the target RHEV hypervisor. To enable the eG agent to obtain the *outside view*, you need to configure the test with the following: ➢ **RHEL MGR HOST** - The IP address/host name of the RHEV manager that the eG agent should connect to ➢ **RHEL MGR PORT** - The port number at which the said RHEV manager listens ➢ **RHEL MGR DOMAIN** - The domain to which the RHEV manager belongs ➢ **RHEL MGR USER** and **RHEL MGR PASSWORD** - The credentials of a user with *read-only access* to the Restful API on the RHEV manager. To know how to create a read-only role and assign it to a user, follow the steps detailed in Section 1.4 of this document. If the RHEV hypervisor being monitored was discovered via an RHEV manager, then the IP address, port number, domain name, and user credentials of the RHEV manager used for discovery will be automatically displayed against the respective parameters. If the RHEV hypervisor being monitored was not discovered via an RHEV manager, but you still want to use an RHEV manager for obtaining the *outside view*, then, you can select any IP address of your choice from the **RHEL MGR HOST** list. By default, this list will be populated with the IP addresses/host names of all the RHEV managers that were configured for the purpose of discovering the RHEV hypervisors. If you select an **RHEL MGR HOST** from this list, then the corresponding port number, domain name, and user credentials will be automatically displayed against the respective parameters. On the other hand, if the RHEV manager that you want to use for metrics collection is not available in the **RHEL MGR HOST** list, then, you can configure an RHEV manager on-the-fly by picking the **Other** option from the **RHEL MGR HOST** list. An **ADD THE RHEV MANAGER DETAILS** window will then pop up. Refer to Section 2.1.1.1 to know how to add an RHEV manager using this window. Once the RHEV manager is added, its IP address, port number, domain name and user credentials will be displayed against the corresponding parameters. 4. **CONFIRM PASSWORD** - Confirm the **RHEL MGR PASSWORD** by retyping it here. 5. **SSL** - If the RHEV manager to which the eG agent should connect is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**.

6. **IGNORE VMS INSIDE VIEW** - Administrators of some high security RHEV environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to **not obtain the 'inside view' of such 'inaccessible' VMs** using the **IGNORE VMS INSIDE VIEW** parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, your **IGNORE VMS INSIDE VIEW** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside view of all VMs on an RHEV host by default.

   Note:

   While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the **IGNORE VMS INSIDE VIEW** text box.

7. **EXCLUDE VMS** - Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the **EXCLUDE VMS** text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your **EXCLUDE VMS** specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the **EXCLUDE VMS** text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.

8. **IGNORE WINNT** - By default, the eG agent does not support the *inside view* for VMs executing on **Windows NT** operating systems. Accordingly, the **IGNORE WINNT** flag is set to **Yes** by default.

9. **INSIDE VIEW USING** - By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the **INSIDE VIEW USING** flag is set to **Remote connection to VM (Windows)**.

   Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the **eG VM Agent** on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs **without domain administrator rights.** Refer to Section 1.4 for more details on the **eG VM Agent**. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the **INSIDE VIEW USING** flag to **eG VM Agent (Windows)**. Once this is done, you can set the **DOMAIN**, **ADMIN USER**, and **ADMIN PASSWORD** parameters to *none*.

10. **DOMAIN, ADMIN USER, ADMIN PASSWORD,** and **CONFIRM PASSWORD** - By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the **DOMAIN** within which the virtual guests reside. The **ADMIN USER** and **ADMIN PASSWORD** will change according to the **DOMAIN** specification. Discussed below are the different values that the **DOMAIN** parameter can take, and how they impact the **ADMIN USER** and **ADMIN PASSWORD** specifications:

    - **If the VMs belong to a single domain** : If the guests belong to a specific domain, then specify the name of that domain against the **DOMAIN** parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the **ADMIN USER** field and the corresponding password in the **ADMIN PASSWORD** field. Confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
    - **If the guests do not belong to any domain (as in the case of Linux guests)**: In this case, specify "none" in the **DOMAIN** field, and specify a local administrator account name in the **ADMIN USER** below.

      Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the **ADMIN USER** against **ADMIN PASSWORD**, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.

      If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the **ADMIN USER** text box, enter the name of the user whose **<USER_HOME_DIR>** (on that Linux guest) contains a **.ssh** directory with the *public key file* named **authorized_keys**. The **ADMIN PASSWORD** in this case will be the *passphrase* of the *public key*; the default *public key file* that is bundled with the eG agent takes the password *eginnovations*. Specify this as the **ADMIN PASSWORD** if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the *passphrase* that you provide while generating the pair. For the detailed procedure on *Implementing Key-based Authentication* refer to Section 2.7 of this document.

	➢ **If the guests belong to different domains -** In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple **DOMAIN** names, multiple **ADMIN USER** names and **ADMIN PASSWORDS** would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the **Click here** hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document. ➢ **If the INSIDE VIEW USING flag is set to 'eG VM Agent (Windows)' -** In this case, the inside view can be obtained without domain administrator privileges. Therefore, set the domain, admin user, and admin password parameters to *none*. 11. **REPORT BY USER** - While monitoring a *RHEV Hypervisor*, the **REPORT BY USER** flag is set to **NO** by default, indicating that by default, the guest operating systems on the hypervisor are identified using the hostname specified in the operating system. On the other hand, while monitoring a *RHEV Hypervisor - VDI*, this flag is set to **YES** by default; this implies that in case of the VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every *username_on_virtualmachinename*. 12. **REPORT POWERED OS** - **This flag becomes relevant only if the REPORT BY USER flag is set to 'Yes'.** If the **REPORT POWERED OS** flag is set to **Yes** (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their *virtual machine name* and not by the *username_on_virtualmachinename*. On the other hand, if the **REPORT POWERED OS** flag is set to **No**, then this test will not report measures for those VMs to which no users are logged in currently.

<table>
<tr><td></td><td colspan="3">13. IGNORESERVICES - Provide a comma-separated list of services that need to be ignored while monitoring. When configuring a service name to exclude, make sure that you specify the Display Name of the service, and not the service Name you see in the Services window on your Windows VM.
14. 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.
15. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each powered-on guest on the RHEV Hypervisor being monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>New automatic services started:
Indicates the number of Windows services with startup type as automatic, which were running in the last measurement period.</td><td>Number</td><td>The detailed diagnosis of this measure lists the services (with startup type as automcatic) that are running.</td></tr>
<tr><td>New automatic services stopped:
Indicates the number of Windows services with startup type as automatic, which were not running in the last measurement period.</td><td>Number</td><td>To know which services stopped, use the detailed diagnosis of this measure (if enabled).</td></tr>
<tr><td></td><td>New manual services started:
Indicates the number of Windows services with startup type as manual, which were running in the last measurement period.</td><td>Number</td><td>Use the detailed diagnosis of this measure to identify the manual services that are running.</td></tr>
</table>

	New manual services stopped: Indicates the number of Windows services with startup type as *manual*, which stopped running in the last measurement period.	Number	To identify the services that stopped, use the detailed diagnosis of this measure.

As stated earlier, by default, clicking on the **Inside View of VMs** layer of a managed *RHEV Hypervisor*, leads you to a page displaying the current status of the virtual guests executing on that server. If you want to override this default setting - i.e., if you prefer to view the tests mapped to the **Inside View of VMs** layer first, and then proceed to focus on individual guest performance, follow the steps given below:

- Edit the **eg_ui.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory
- Set the **LAYERMODEL_LINK_TO_VIRTUAL** flag in the file to **false**; this is set to **true** by default.
- Save the **eg_ui.ini** file.

Doing so ensures that as soon as the **Inside View of VMs** layer is clicked, the list of tests mapped to that layer appears, as depicted by Figure 2.24.

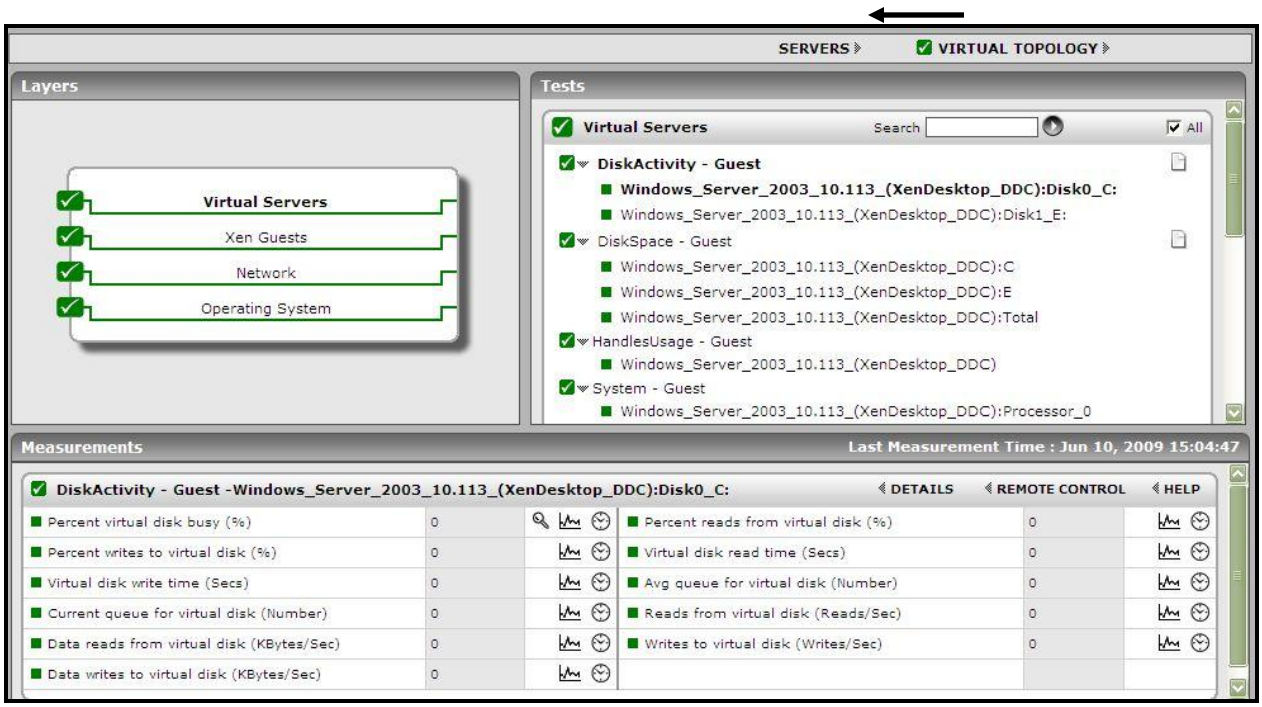

Figure 2.24: The tests mapped to the Virtual Servers layer

If you now want the **Server view** of Figure 2.12, simply click on the **SERVERS** link above the list of tests in Figure 2.24 (indicated by the arrow).

Clicking on any of the guests in the **Server view** leads you to Figure 2.25 that displays all the performance metrics extracted from that guest, in real-time. You are thus enabled to cross-correlate across the various metrics, and quickly detect the root-cause of current/probable disturbances to the internal health of a guest. To view the time-of-day variations in a measure, you can view its graph by clicking on that measure in Figure 2.25.

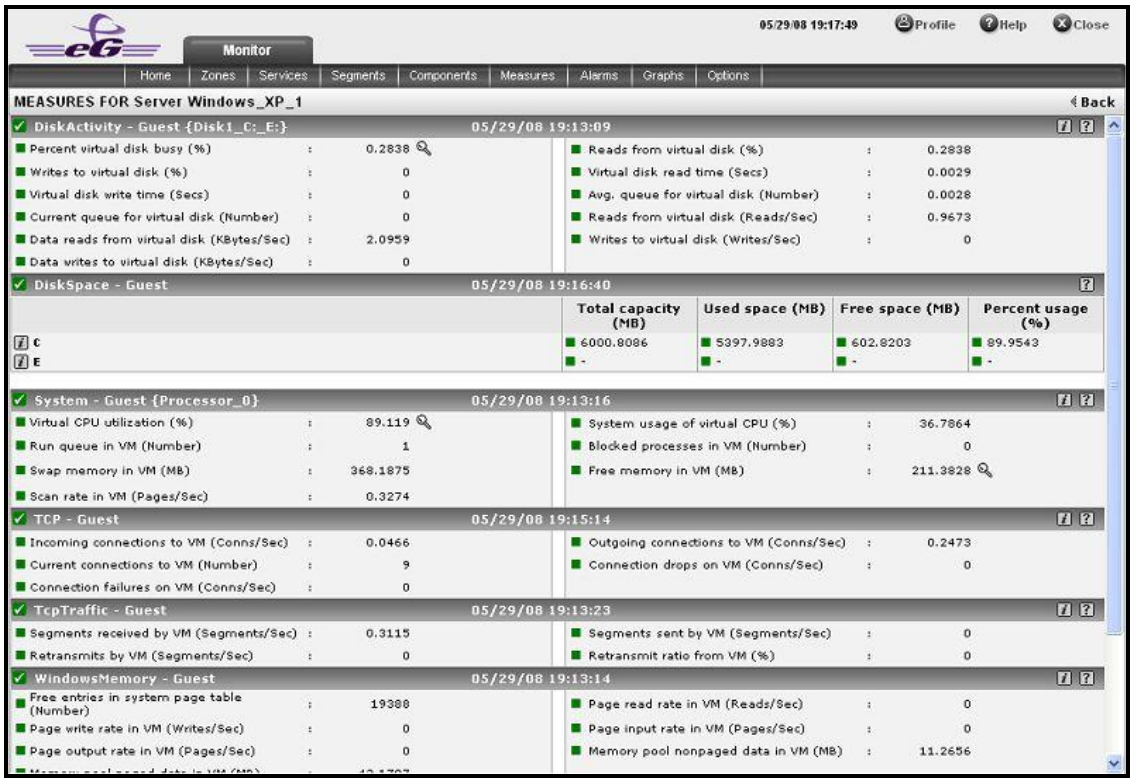

Figure 2.25: Measures pertaining to a chosen guest

To view real-time graphs of pre-configured measures (pertaining to the RHEV Hypervisor and the guests operating on it), click on the **LIVE GRAPH** link in Figure 2.12. Alternatively, you can click on the icon that appears in the **Tests** panel of the layer model page when the **Inside View of VMs** layer is clicked to view the live graph. The graph display that appears subsequently has been organized in such a way that next to every host-pertinent measure graph, the closely related guest-specific measure graph appears. This way, you can easily compare and correlate how well the physical CPU resources are being utilized by both the RHEV Hypervisor and the guests. On the basis of this analysis, you can proactively isolate potential performance issues, and also determine the root-cause of the issue - is it the RHEV Hypervisor? or is it the virtual guest? If you access this page from the **LIVE GRAPH** link in Figure 2.12, then, by default, you will view live graphs pertaining to the *RHEV Hypervisor*. However, you can select a different virtualized component-type and a different virtualized component using the **type** and **Component Name** lists (respectively).

2.7 Troubleshooting

2.7.1 Troubleshooting the Failure of the eG Remote Agent to Connect to or Report Measures for Linux Guests

By default, the eG agent uses secure shell (SSH) to connect to Linux guests, and collect performance metrics from them. Password Authentication is the default method for SSH connections in eG Enterprise. If the eG agent fails to report measures for a Linux guest or is unable to connect to a guest, it could imply that the Linux guest does not support SSH or that password authentication is not supported by the SSH daemon running on the Linux guest. Under such circumstances, you can perform either of the following:

- Enable Password Authentication in the SSH daemon on the Linux guest; or,
- Implement Key-Based Authentication between eG agent and the SSH daemon of the Linux guest.

If you pick option (1), then follow the steps given below to enable password authentication:

- Login to the Linux guest to be monitored.
- Edit the **sshd_config** file in the **/etc/ssh** directory.
- Check whether the **Password Authentication** flag in the **sshd_config** file is set to **no**. If so, set it to **yes**.
- Then, save the file and restart/signal the SSH daemon (eg., using **kill -1 <sshd_config PID>**).

On the contrary, if you choose to enable key-based authentication [i.e, option (2)], then you will have to generate a public/private key pair. A public/private key pair is available in the **<EG_INSTALL_DIR>\agent\sshkeys** directory (on Windows; on Unix, this will be **/opt/egurkha/agent/sshkeys**) of the eG agent. While the private key is available in the file named **id_rsa**, the public key is contained within the file **authorized_keys**. You now have the option to proceed with the default keys or generate a different key pair. If you decide to go with the keys bundled with the eG agent, do the following:

- To enable key-based authentication, the private key should remain in the **<EG_INSTALL_DIR>\agent\sshkeys** directory (on Windows; on Unix, this will be **/opt/egurkha/agent/sshkeys**), and the public key should be copied to each of the Linux guests to be monitored. To achieve this, first login to the Linux guest to be monitored as the eG user.
- Create a directory named **.ssh** in the **<USER_HOME_DIR>** on the guest operating system, using the command: **mkdir ~/.ssh**.
- Next, copy the **authorized_keys** file from the **<EG_INSTALL_DIR>\agent\sshkeys** directory (on Windows; on Unix, this will be **/opt/egurkha/agent/sshkeys**) on the eG remote agent host to the **<USER_HOME_DIR>/.ssh** directory on the Linux guest.
- Make sure that the permission of the **.ssh** directory and the **authorized_keys** file is **700**.
- Finally, on the eG manager host, edit the **<EG_INSTALL_DIR>\manager\config\eg_tests.ini** file. Against the **EgJavaSSHKeyFile** parameter, enter: **agent/sshkeys/id_rsa.pub**, and save the file.

On the other hand, if you want to generate a new key pair, then do the following:

- Login to any Linux host in your environment (even a Linux VM) as an eG user.
- From the **<USER_HOME_DIR>**, execute the command: **ssh-keygen -t rsa**. Upon executing the command, you will be requested to specify the full path to the file to which the key is to be saved. By default, a directory named **.ssh** will be created in the **<USER_HOME_DIR>**, to which the key pair will be saved. To go with the default location, simply press **Enter**.

```
Generating public/private rsa key pair.
Enter file in which to save the key (/home/egurkha/.ssh/id_rsa):
```

- Next, you will be prompted to provide a pass phrase. Provide any pass phrase of your choice.

```
Enter passphrase (empty for no passphrase): eginnovations
Enter same passphrase again: eginnovations
```

- If the key pair is created successfully, then the following messages will appear:

```
Your identification has been saved in /hom
e/egurkha/.ssh/id_rsa.
Your public key has been saved in /home/egurkha/.ssh/id_rsa.pub.
The key fingerprint is:
09:f4:02:3f:7d:00:4a:b4:6d:b9:2f:c1:cb:cf:0e:e1 dclements@sde4.freshwater.com
```

- The messages indicate that the private key has been saved to a file named **id_rsa** in the **<USER_HOME_DIR>/.ssh**, and the public key has been saved to a file named **id_rsa.pub** in the same directory. Now, to enable key-based authentication, login to the Linux guest to be monitored as the eG user.
- Create a directory named **.ssh** in the **<USER_HOME_DIR>** on the guest operating system, using the command: **mkdir ~/.ssh**.
- Next, copy the **id_rsa.pub** file from the **<USER_HOME_DIR>/.ssh** directory on the Linux host to the **<USER_HOME_DIR>/.ssh** directory on the Linux guest.
- Ensure that the **id_rsa.pub** file on the Linux guest is renamed as **authorized_keys**.
- Repeat this procedure on every Linux guest to be monitored.
- Then, lock the file permissions down to prevent other users from being able to read the key pair data, using the following commands:

  chmod go-w ~/
  chmod 700 ~/.ssh
  chmod go-rwx ~/.ssh/*

- Finally, on the eG manager host, edit the **<EG_INSTALL_DIR>\manager\config\eg_tests.ini** file. Against the **EgJavaSSHKeyFile** parameter, enter: **agent/sshkeys/id_rsa.pub**, and save the file.

Instead of choosing between the authentication modes (Password or Key-based), you can also disable the usage of the Java SSH client, and use **plink** to connect to Linux guests. To achieve this, follow the steps given below:

- Edit the **eg_tests.ini** file in the **/opt/egurkha/manager/config** directory (on Unix; on Windows, this will be **<EG_INSTALL_DIR>\manager\config** directory).
- Set the **JavaSSHForVm** flag in the **[AGENT_SETTINGS]** section of the file to **false**; by default, this is set to **true** indicating that the eG agent uses Java SSH by default. By setting the flag to **false**, you can ensure that the eG agent does not use Java SSH, and instead uses the **plink** command to connect to Linux guests.
- The **plink** command exists in the **<EG_INSTALL_DIR>\lib\vmgfiles** directory (on Windows; on Unix, this will be **/opt/egurkha/lib/vmgfiles**) of the eG agent. To use the **plink** command, you need to explicitly configure the SSH keys, so that the eG agent is able to communicate with the Linux guests using SSH. To do this, follow the steps given below:
  - Go to the command prompt and switch to the directory containing the **plink** command.
  - Then, execute the **plink** command to connect to any of the Linux-based virtual machines on the RHEV server. The syntax for the **plink** command is as follows:

    plink -ssh <user>@<IP_of_target_host> <command>

    For example, assume that you want to connect to the virtual machine, **192.168.10.7**, as user **john** with password **john**, to know its hostname. The syntax of the **plink** command in this case will be:

plink -ssh john@192.168.10.7 hostname, where **hostname** is the command to be executed on the remote host for extracting its hostname.

- To ensure that you do not connect to an imposter host, **SSH2.x** presents you with a unique host key fingerprint for that host, and requests your confirmation to save the displayed host key to the cache.

```
The server's host key is not cached in the registry. You have no guarantee
that the server is the computer you think it is.
The server's rsa2 key fingerprint is:<host key>
If you trust this host, enter "y" to add the key to PuTTY's cache and carry
on connecting.
If you want to carry on connecting just once, without adding the key to the
cache, enter "n".
If you do not trust this host, press Return to abandon the connection.
Store key in cache? (y/n) y
```

Once you confirm the host key storage and provide the user's password to connect to the virtual guest, this message will not appear during your subsequent attempts to connect to any Linux-based virtual machine on the monitored RHEV server. In other words, the eG agent will be able to execute tests on all Linux guests on the target RHEV server without any interruption. Therefore, press **y** to confirm key storage.

Chapter 3

The RHEV Hypervisor - VDI Monitoring Model

In some environments, the virtual guests hosted on RHEV servers may be used to support desktop applications. Administrators of such virtual environments would want to know the following:

- How many desktops are powered on simultaneously on the RHEV server?
- Which users are logged on and when did each user login?
- How much CPU, memory, disk and network resources is each desktop taking?
- What applications are running on each desktop?
- Which RHEV server is a virtual guest running on?
- When was a guest moved from an RHEV server? Which RHEV server was the guest moved to?
- Why was the guest migrated? What activities on the RHEV server caused the migration?

Using the *RHEV Hypervisor - VDI* model (see Figure 3.1), administrators can find quick and accurate answers to all the queries above, and also receive a complete 'desktop view', which allows them to get up, close with the performance of every guest OS hosted by the RHEV server and detect anomalies (if any) in its functioning.

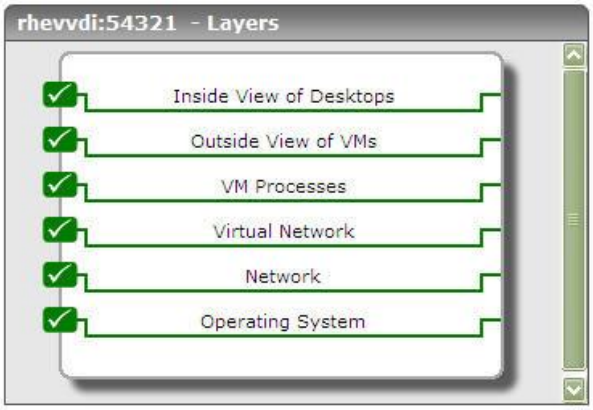

Figure 3.1: The RHEV Hypervisor - VDI Monitoring Model

The tests mapped to the bottom five layers of Figure 3.1 have already been discussed in the previous chapter. The **Inside View of Desktops** layer runs the same tests as the **Inside View of VMs** layer of the *RHEV Hypervisor* model. However, the difference lies in what these 'inside view' tests monitor; in the case of the VDI environment these tests report metrics for the user currently logged into the VM - not the VM itself as in the case of the *RHEV Hypervisor* model.

Chapter 4

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 **RHEV**. 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.