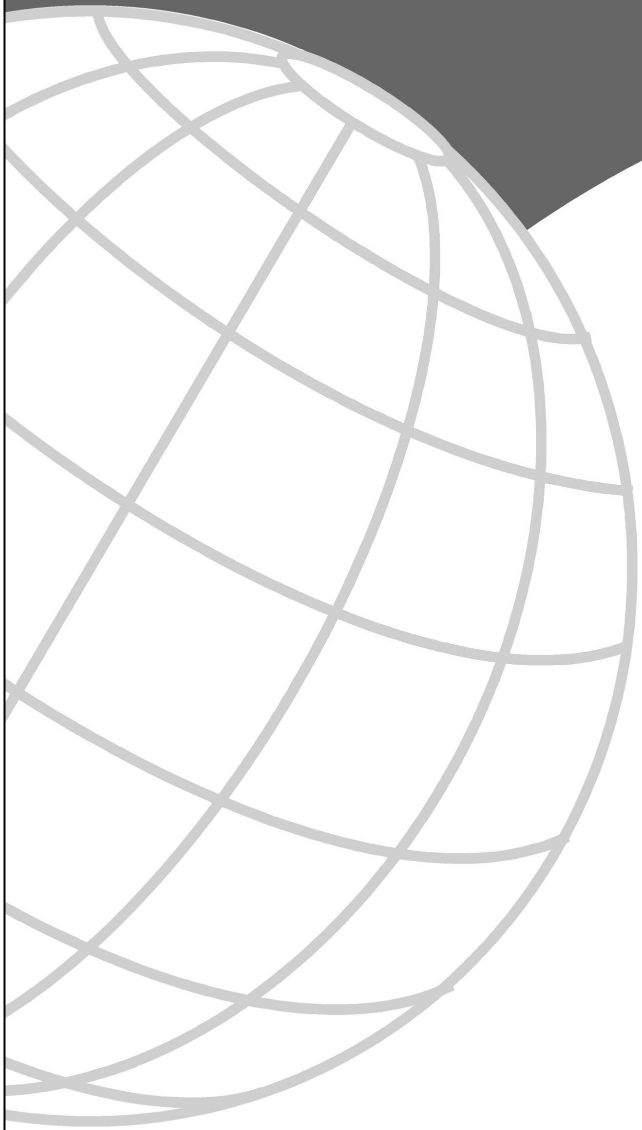

Enabling Service Excellence

Monitoring Citrix XenServer

eG Enterprise v6.1.2

Restricted Rights Legend

The information contained in this document is confidential and subject to change without notice. No part of this document may be reproduced or disclosed to others without the prior permission of eG Innovations Inc. eG Innovations Inc. makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

Trademarks

Microsoft Windows, Windows 2008, Windows 2012, Windows 2016, Windows 7, Windows 8 and Windows 2010 are either registered trademarks or trademarks of Microsoft Corporation in United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Copyright

Table of contents

Table of Figures

1

Introduction

Citrix XenServer, a member of the Citrix Delivery Center product family, allows the rapid deployment of high-performance virtual machines, and the management of their related storage and networking resources from a single management console.

Xen provides fast, secure, open source virtualization that allows multiple operating system instances to run as *Xen Virtual Machines* or *XenVMs* on a single physical x86 computer. Xen is based on a para-virtualization technology, which presents an abstraction of the hardware that is similar but not identical to the underlying physical hardware. Para-virtualization techniques require modifications to the guest operating systems that are running on the VMs. As a result, the guest operating systems are aware that they are executing on a VM - allowing for near-native performance.

Figure 1.1 shows the architecture of Xen server hosting four VMs (Domain 0, VM 1, VM 2, and VM 3). This architecture includes the Xen Virtual Machine Monitor (VMM) or the Xen Hypervisor, which abstracts the underlying physical hardware and provides hardware access for the different virtual machines. Since the Xen hypervisor is installed on the ‘bare metal’, it lies between the physical hardware and the base operating system, and effectively decouples the guest operating systems and the applications executing within from the underlying physical server. Figure 1 shows the special role of the VM called Domain 0 or the control domain. Only the control domain can access the control interface of the VMM, through which other VMs can be created, destroyed, and managed. Management and control software runs in the control domain. Administrators can create virtual machines with special privileges—such as VM 1—that can directly access the hardware through secure interfaces provided by Xen. Administrators can create other virtual machines that can access the physical resources provided by Domain 0’s control and management interface in Xen.

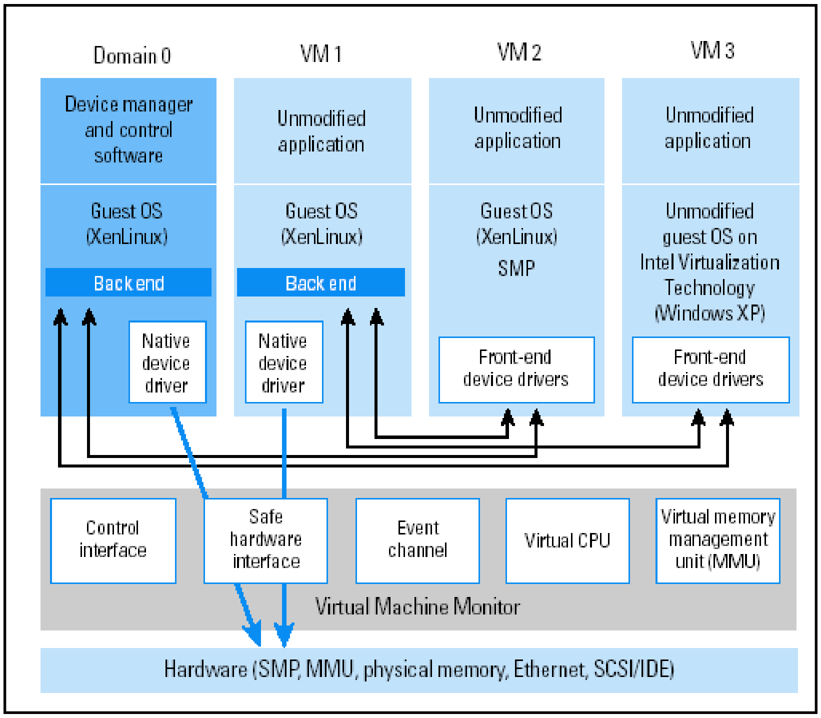

Figure 1.1: The architecture of a Citrix XenServer

In this example, the guest operating systems in VM 1 and in VM 2 are modified to run above Xen and also have Xen-aware drivers to enable high performance. Near-native performance can be achieved through this approach. Unmodified guest operating systems are also supported using new processor extensions from Intel (VT) and AMD (AMD-V).

1.1 XenServer Deployment Models

Xen virtualization is being deployed in a couple of different ways. A single physical XenServer can be host to multiple server applications. For instance, a Citrix server, a Web server, and an Oracle database server can all be operated on three VMs on the same physical server.

Alternatively, virtual desktops (instead of servers) can be published on the XenServer, and users can connect to virtual desktops to run their client applications. In this scenario, each user is assigned to a virtual desktop and can run his/her applications within that desktop.

1.2 Challenges in Monitoring XenServers

Xen enables administrators to implement two types of virtualization: para-virtualized VMs, which can enhance performance but require guest OS modifications, and fully virtualized VMs, which are highly portable and do not require guest OS modifications. With more and more IT environments these days turning to virtualization for easily and effectively managing their systems and resources, the flexibility that Xen offers in the implementation of virtualization makes it a coveted solution, especially for large infrastructures. Such infrastructures would typically have tens of virtual machines configured on a single XenServer host, with all VMs and applications sharing the physical CPU/memory/storage resources of the base XenServer host. In such infrastructures therefore, a resource contention at the host is bound to impact the resource allocation to the guests, thereby affecting the performance of the applications executing on the guests. Likewise, resource-intensive applications running on a VM can also drain the VM of the physical resources allocated to it. Therefore, whenever an application executing on a VM experiences a slowdown, administrators often take hours to determine where the root-cause of the problem lies - with the XenServer host or with the guest VM?

Moreover, the different virtualization deployment models pose different 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 compounds the management troubles of administrators. In addition, while in-depth monitoring of each of the applications is important in the server application virtualization approach, monitoring of the desktop need not be as in-depth. 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.

1.3 How eG Enterprise Monitors Citrix XenServers?

eG Enterprise provides two specialized monitoring models for XenServers - a Citrix XenServer (see Figure 2) model for monitoring XenServers that are hosting server applications, and a Citrix XenServer - VDI model that focuses on the performance of the virtual desktops executing on physical XenServers.

Both these models have been discussed elaborately later in this document.

Regardless of the model used, eG Enterprise adopts a patented ‘In-N-Out’ approach to monitoring Xen virtualized environments. In this approach, administrators have the option to choose between agent-based and agentless methodologies for monitoring the virtualized environment end-to-end, and determining the following:

- The resource usage of the physical server - i.e., the XenServer host
- The resource usage of each of the guests in relation to the physical resources available at the host-level (i.e., the “outside” view)
- The usage of allocated resources by the guests (i.e., the “inside” view)

To retrieve these useful statistics, eG Enterprise offers the deployment models discussed below.

1.3.1 Agent Deployment Models

As mentioned earlier, administrators can choose between the agent-based and agentless approaches to Xen monitoring. In case of the “agent-based” approach, a single eG agent deployed on the control domain (i.e., Domain 0) of the XenServer can be configured to monitor the host and each of the guest operating systems executing on it (see Figure 1.2). The control domain of the XenServer runs a Linux variant. Hence, the eG agent for Linux should be deployed on the XenServer. When installing the eG agent, make sure that you have at least 350MB of free disk space on Domain 0. The procedure for deploying the eG agent on the control domain is exactly the same as the process for installing the eG agent on Linux, which is described in the *eG Installation Guide*.

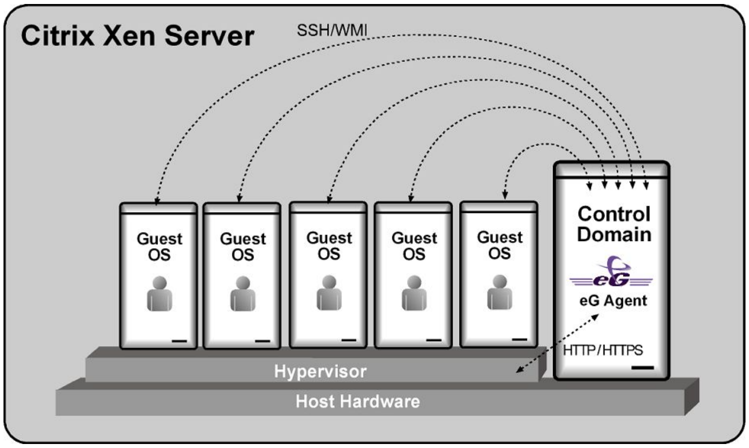

Figure 1.2: Monitoring a Citrix XenServer in an agent-based manner

In case of the “agentless” approach, an eG agent deployed on a remote Windows or Linux host in the environment can be configured to monitor the loading of the physical XenServer, the relative loading of each of the guests on the server, and how well each guest is using the resources allocated to it (see Figure 1.3). Refer to the *eG Installation Guide* for a detailed procedure on how to install and configure an eG agent on Windows/Linux.

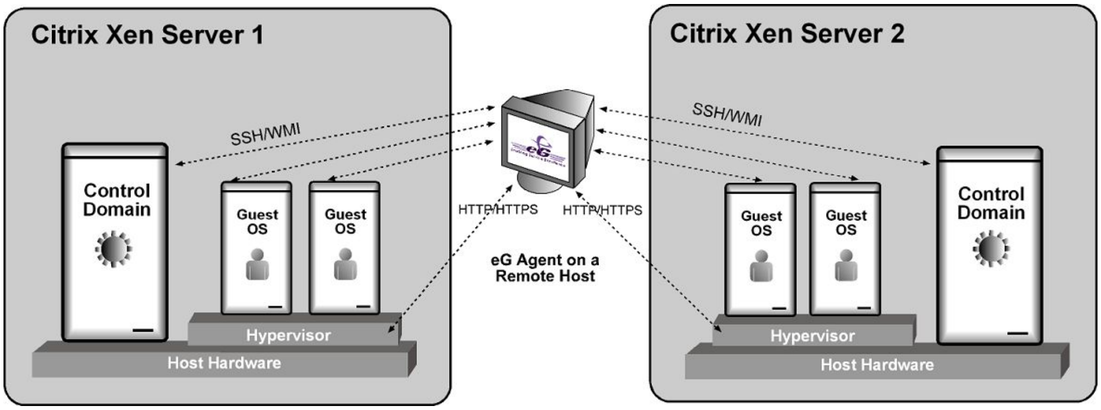

Figure 1.3: Monitoring Citrix XenServers in an agentless manner

Figure 1.2 and Figure 1.3 clearly indicate that the eG agent (whether on a control domain or on a remote Windows/Linux host) remotely communicates with each guest via SSH/WMI (depending upon the virtual OS in use) to obtain the “inside view” of the guests. 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 to collect “inside view” metrics from the VMs **without domain administrator rights.** Refer to Section Section **1.4** for more details on the **eG VM Agent**.

1.3.2 Pre-requisites for Monitoring a XenServer

There are several pre-requisites for an eG agent to be able to monitor a XenServer and the guest VMs hosted on it.

1. **If an eG agent is installed on the control domain, allow the eG agent to communicate back to the eG management console:** Make sure that the firewall on the XenServer is configured to allow outbound traffic from the eG agent to the eG management console. The port used for this communication is determined at the time the eG manager and agents are installed in your environment; port 7077 is the default. To configure the agent-manager communication, do the following:

   - Login to the XenServer host.
   - Edit the **iptables** file in the **/etc/sysconfig/** directory.
   - To open the eG manager port, insert the following line anywhere in the file, but before the **REJECT** line:
     -A RH-Firewall-1-INPUT -p tcp -m tcp -dport 7077 -j ACCEPT
   - Save the file.
   - Restart the network service by issuing the command: **/sbin/service iptables restart**

2. **Enable auto-discovery of VMs by eG Enterprise:** Xentools must be installed on all guest operating systems hosted on a XenServer. Using Xentools, the eG agent determines the IP addresses of the guest VMs and the operating systems that they are configured with. If the eG Enterprise monitor shows “N/A” against the IP address field or the operating system type of any VM, this is usually a good indicator that Xentools has not been installed on that VM.

3. **Enable the eG agent to access the XenServer API (for agentless monitoring or for agent-based monitoring)**:

   - In order to ensure that the eG agent uses XenServer API to discover the guest operating systems executing on a target XenServer host, all the tests that the agent executes should be configured with the name and password of a registered user of the XenServer.
   - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag of all tests executed by the eG agent is set to **No** by default.
     If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes** for all tests executed by the eG agent, so that the eG agent communicates with the XenServer using HTTPS.

     Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

- Obtain the server-certificate for the XenServer
- Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

4. **Enabling the eG agent to collect "inside view" metrics from Linux guests**:

   To allow the eG agent to obtain the "inside view" of Linux VMs, simply ensure that SSH is enabled on all Linux guests to be monitored.

5. **Enabling the eG agent to collect "inside view" metrics from Windows guests, without using the eG VM Agent**:

   To allow the eG agent to obtain the "inside view" of Windows VMs without using the **eG VM Agent**, the following pre-requisites need to be fulfilled:

   - 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 1.3.4 for a step-by-step procedure to achieve this.
   - In case of VMs with the Windows XP/Windows 2003/Windows Vista operating systems, the firewall on the guest should be explicitly configured to allow Windows File and Print Sharing services which are required for the eG agent on the Xen host to communicate with the guest operating system.
   - Make sure that the XenServer firewall allows the eG agent to communicate with the Windows File and Print Sharing port (typically, this is port 139). To configure this communication, do the following on the XenServer host:
     - Login to the XenServer host.
     - ***Edit the*** **iptables** file in the **/etc/sysconfig/** directory.
     - To open the Windows File and Print Sharing port, insert the following line anywhere in the file, but before the **REJECT**line:

       -A RH-Firewall-1-INPUT -p tcp -m tcp -dport 139 -j ACCEPT
     - Save the file.
     - Restart the network service by issuing the command: **/sbin/service iptables restart**
   - Also, ensure that the XenServer firewall allows the eG agent on the server to communicate with the Linux guests of the server using SSH. For instance, to allow the eG agent to communicate with the Linux guests listening on **port 22** on a XenServer, do the following:
     - Login to the XenServer host.
     - ***Edit the*** **iptables** file in the **/etc/sysconfig/** directory.
     - To open the SSH port 22, insert the following line anywhere in the file, but before the **REJECT**line:

       -A RH-Firewall-1-INPUT -p tcp -m tcp -dport 22 -j ACCEPT

- Save the file.
- Restart the network service by issuing the command: **/sbin/service iptables restart**

➢ To enable the eG agent to communicate with the guest operating systems, an administrative account login and password (either a local account or a domain account) must be provided when configuring the eG monitoring capabilities; different logins can be provided for different VMs on the same XenServer. However, ensure that this account is available or is explicitly created on each of the virtual machines being monitored on a XenServer.

➢ Set the **INSIDE VIEW USING** flag for all the "inside view" tests to **Remote connection to VM (Windows)**.

6. **Enabling the eG agent to collect "inside view" metrics from Windows guests, with the help of the eG VM Agent**:

   To allow the eG agent to obtain the "inside view" of Windows VMs using the **eG VM Agent**, the following pre-requisites need to be fulfilled:

   ➢ Install the eG VM Agent

   ➢ Set the **INSIDE VIEW USING**flag for all the "inside view" tests to **eG VM Agent (Windows)**.

7. **Ensure connectivity from the eG agent to the XenServer being monitored and the VMs:** Since the same agent is used to monitor the outside view of the VMs and the inside view of the VMs, ensure that the agent has IP connectivity to the XenServer and to at least one of the network interfaces of the VMs.

8. **If agentless monitoring is used, ensure communication between the eG remote agent (which performs the agentless monitoring) and the individual VMs**.

   ➢ For monitoring a Windows VM, TCP port 139 must be accessible from the remote agent to the VM.

   ➢ To enable the remote agent (on Windows) to obtain the inside view of Windows VMs, the **eGurkhaAgent** service should run using *domain administrator* privileges.

   ➢ 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.

9. **Ensure that the XenServer is configured to allow remote performance monitoring**:

   When configuring monitoring for XenServers in the eG Enterprise administration console, you will be prompted to enter a user name which will be used by the agent to collect performance metrics from the XenServer. For monitoring XenServer 5.5 (or below), you must specify the "root" user credentials for the eG agent to be able to collect metrics. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user** when monitoring **XenServer 5.6 (or above)**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   ➢ Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the

Component type, and click the **Modify** button corresponding to the target XenServer.

- In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the os is set to **Xen** and the **Mode** is set to **SSH**.
- Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via **SSH**.
- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **xen user** with *Read-only* privileges.

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

1.3.4 Enabling ADMIN$ Share Access on Windows Virtual Guests

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 1.4 appears indicating the same.

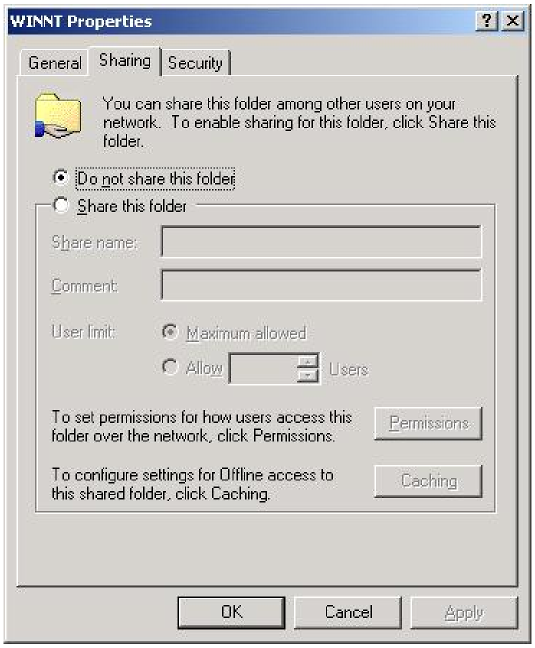

You can share this folder among other users on your network. To enable sharing for this folder, click Share this folder.

● Do not share this folder

○ Share this folder

● Maximum allowed

○ Allow

To set permissions for how users access this folder over the network, click Permissions.

To configure settings for Offline access to this shared folder, click Caching.

Figure 1.4: The ADMIN$ share does not exist

On the other hand, if the **ADMIN$** share pre-exists, Figure 1.5 appears. In such a case, first, remove the **ADMIN$** share by selecting the **Do not share this folder** option from Figure 1.5 and clicking the **Apply** and **OK** buttons. After this, you will have to repeat step 1 of this procedure to open Figure 1.4. Then, proceed as indicated by step 3 onwards.

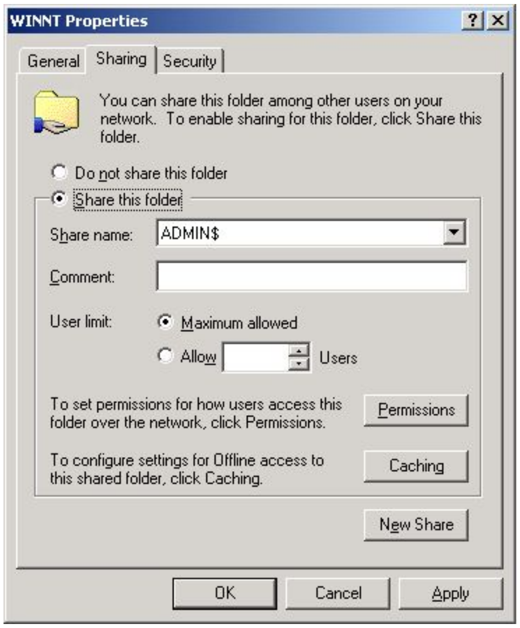

WINNT Properties

You can share this folder among other users on your network. To enable sharing for this folder, click Share this folder.

○ Do not share this folder
● Share this folder

Share name: ADMIN$

Comment:

User limit: ● Maximum allowed
○ Allow Users

To set permissions for how users access this folder over the network, click Permissions.

To configure settings for Offline access to this shared folder, click Caching.

Figure 1.5: Admin$ share pre-exists

3. To create (or re-create) the **ADMIN$** share, select the **Share this folder** option from Figure 1.6, and provide **ADMIN$** share against the **Share name** text box (see Figure 1.6).

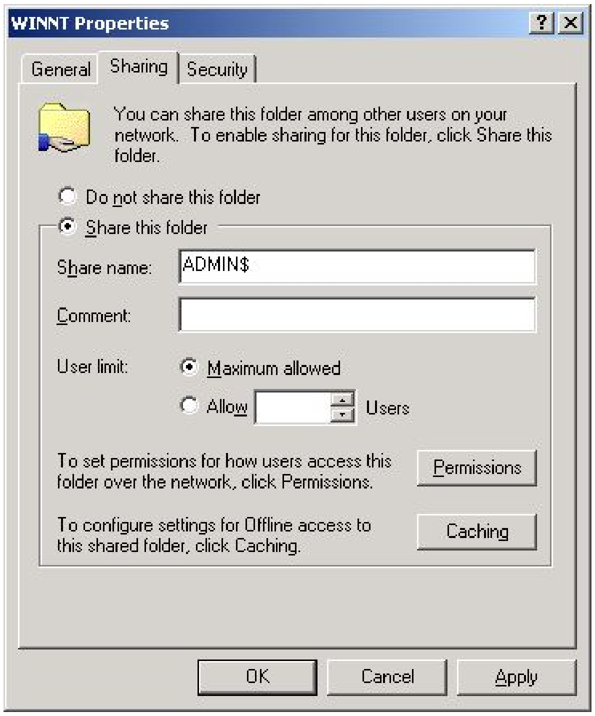

WINNT Properties

You can share this folder among other users on your network. To enable sharing for this folder, click Share this folder.

○ Do not share this folder
● Share this folder

Share name: ADMIN$

Comment:

User limit: ● Maximum allowed
○ Allow Users

To set permissions for how users access this folder over the network, click Permissions.

To configure settings for Offline access to this shared folder, click Caching.

Figure 1.6: 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 XenServer tests. To grant the access permissions, click on the **Permissions** button in Figure 1.6.

5. By default, the **ADMIN$** share can be accessed by **Everyone** (see Figure 1.7). To grant access rights to a specific administrative (local/domain) user, select the **Add** button in Figure 1.7. When Figure 1.8 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.

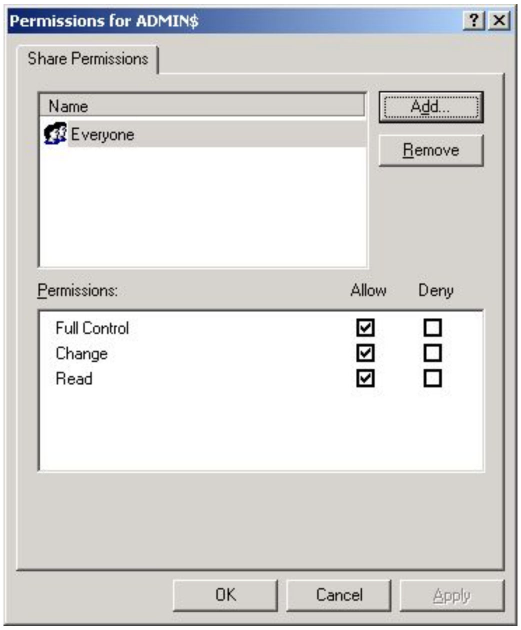

Permissions:	Allow	Deny
Full Control	☑	□
Change	☑	□
Read	☑	□

Figure 1.7: Clicking the Add button

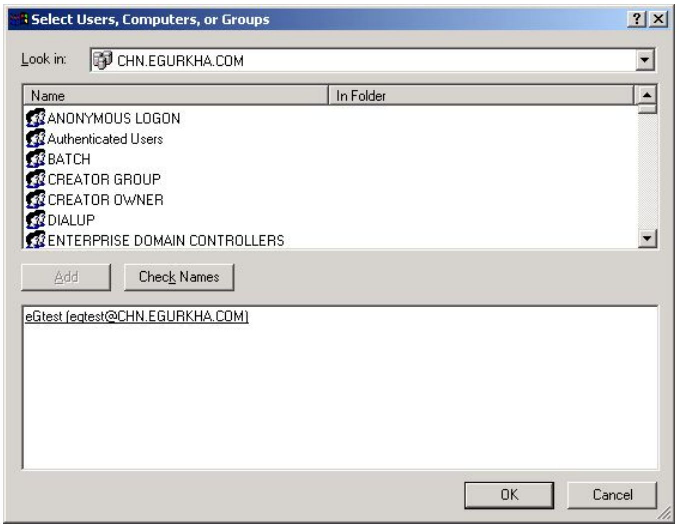

Look in: CHN.EGURKHA.COM

Name	In Folder
ANONYMOUS LOGON	
Authenticated Users	
BATCH	
CREATOR GROUP	
CREATOR OWNER	
DIALUP	
ENTERPRISE DOMAIN CONTROLLERS	

eGtest (egtest@CHN.EGURKHA.COM)

Figure 1.8: Selecting the administrative user to whom access rights are to be granted

6. Finally, click the **OK** button. You will then return to Figure 1.7, where the newly added administrator account will appear (see Figure 1.9).

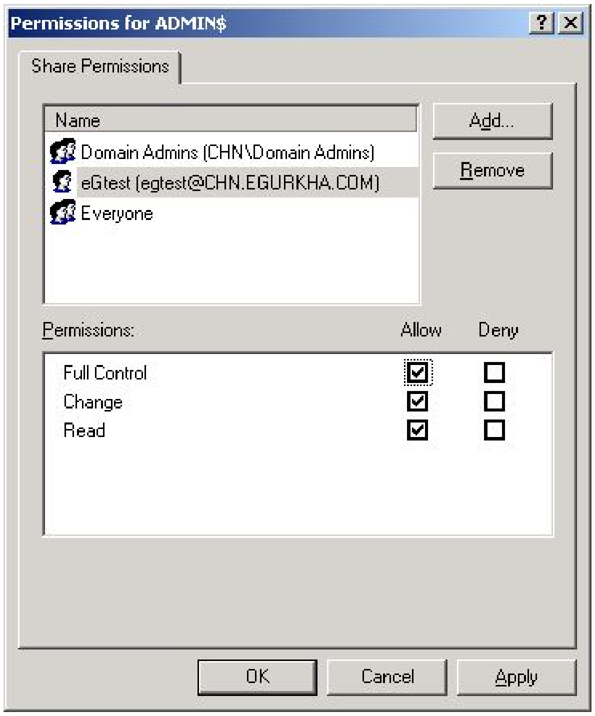

Name
- Domain Admins (CHN\Domain Admins)
- eGtest (egtest@CHN.EGURKHA.COM)
- Everyone

Permissions:	Allow	Deny
Full Control	☑	□
Change	☑	□
Read	☑	□

Figure 1.9: The administrator account granted access permissions

7. Select the newly added administrator account from Figure 1.9, 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 1.9 to register the changes.

9. Once you return to Figure 1.6, click on the **Security** tab (see Figure 1.10) to define the security settings for

the **ADMIN$** share.

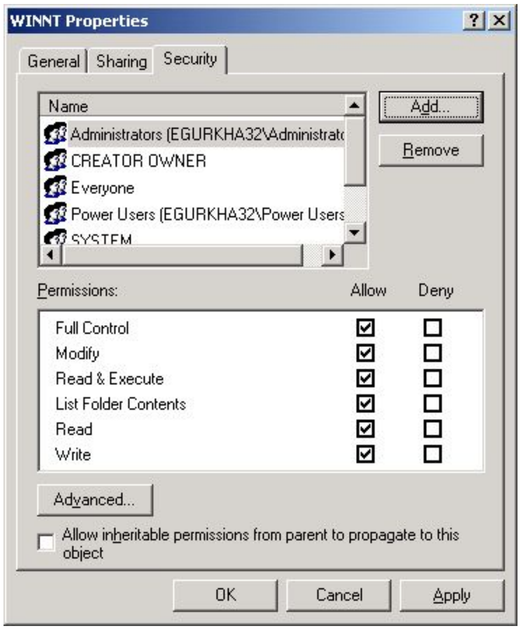

Name
Administrators (EGURKHA32\Administrato
CREATOR OWNER
Everyone
Power Users (EGURKHA32\Power Users
SYSTEM

Permissions:	Allow	Deny
Full Control	☑	□
Modify	☑	□
Read & Execute	☑	□
List Folder Contents	☑	□
Read	☑	□
Write	☑	□

□ Allow inheritable permissions from parent to propagate to this object

Figure 1.10: 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 1.10, pick a domain from the **Look in** list of Figure 1.11, select the said administrator account from the domain users list below, and click the **Add** button (in Figure 1.11) to add the chosen account. Then, click the **OK** button in Figure 1.11.

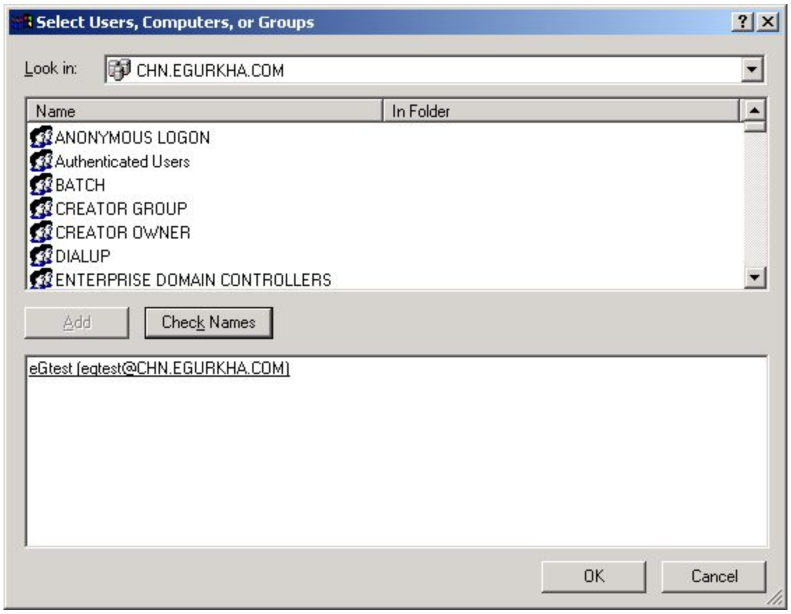

Look in: CHN.EGURKHA.COM

Name	In Folder
ANONYMOUS LOGON	
Authenticated Users	
BATCH	
CREATOR GROUP	
CREATOR OWNER	
DIALUP	
ENTERPRISE DOMAIN CONTROLLERS	

eGtest (egtest@CHN.EGURKHA.COM)

Figure 1.11: Adding the administrator account

11. This will bring you back to the **Security** tab, but this time, the newly added domain administrator account

will be listed therein as indicated by Figure 1.12.

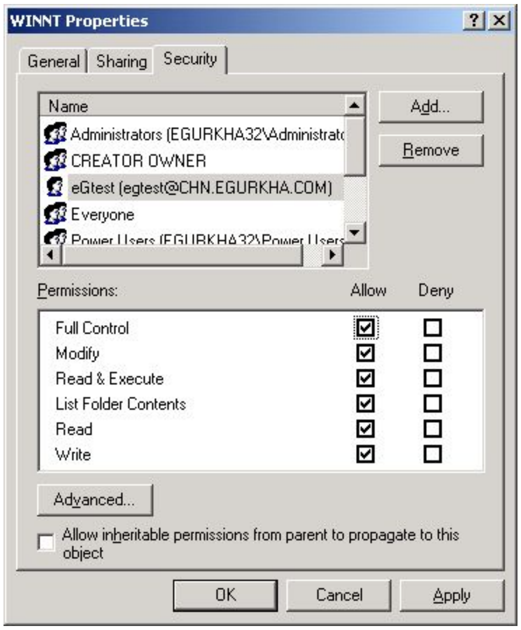

Name
Administrators (EGURKHA32\Administrato
CREATOR OWNER
eGtest (egtest@CHN.EGURKHA.COM)
Everyone
Power Users (EGURKHA32\Power Users

Permissions:	Allow	Deny
Full Control	☑	□
Modify	☑	□
Read & Execute	☑	□
List Folder Contents	☑	□
Read	☑	□
Write	☑	□

□ Allow inheritable permissions from parent to propagate to this object

Figure 1.12: The Administrator account in the Security list

12. Finally, click the **Apply** and **OK** buttons in Figure 1.12.

1.4 Configuring Windows Virtual Machines to Support the eG Agent's Inside View Using the eG VM Agent

To provide the inside view of a Unix 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 VM 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**.

1.4.1 Installing the eG VM Agent

- Manually install the eG VM Agent on every Windows VM using the executable that eG Enterprise provides;

- 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;
- Connect to each Windows VM and silently install the eG VM Agent on it, without using the executable that eG Enterprise provides.

The first and fourth installation options alone are discussed here.

1.4.1.1 Using the Executable Provided by eG Enterprise

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 1.13 then appears. Click on the **Next** button in Figure 1.13 to continue.

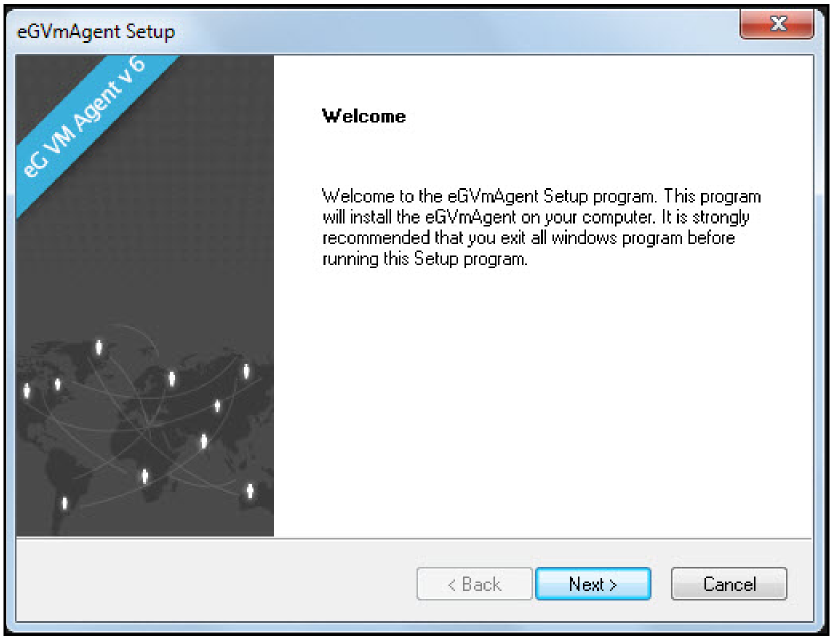

Welcome

Welcome to the eGVmAgent Setup program. This program will install the eGVmAgent on your computer. It is strongly recommended that you exit all windows program before running this Setup program.

Figure 1.13: Welcome screen of the eG VM Agent installation wizard

3. When Figure 1.14 appears, click on **Yes** to accept the displayed license agreement.

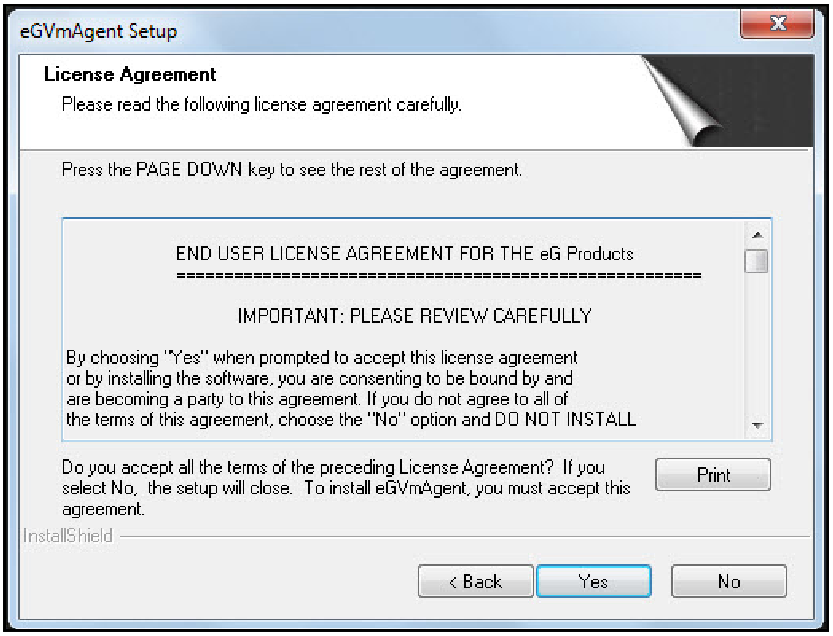

eGVmAgent Setup

License Agreement

Please read the following license agreement carefully.

Press the PAGE DOWN key to see the rest of the agreement.

END USER LICENSE AGREEMENT FOR THE eG Products
===

IMPORTANT: PLEASE REVIEW CAREFULLY

By choosing "Yes" when prompted to accept this license agreement
or by installing the software, you are consenting to be bound by and
are becoming a party to this agreement. If you do not agree to all of
the terms of this agreement, choose the "No" option and DO NOT INSTALL

Do you accept all the terms of the preceding License Agreement? If you select No, the setup will close. To install eGVmAgent, you must accept this agreement.

Figure 1.14: Accepting the license agreement

4. Use the **Browse** button in Figure 1.15 to indicate the location in which the agent should be installed, and click the **Next** button to proceed.

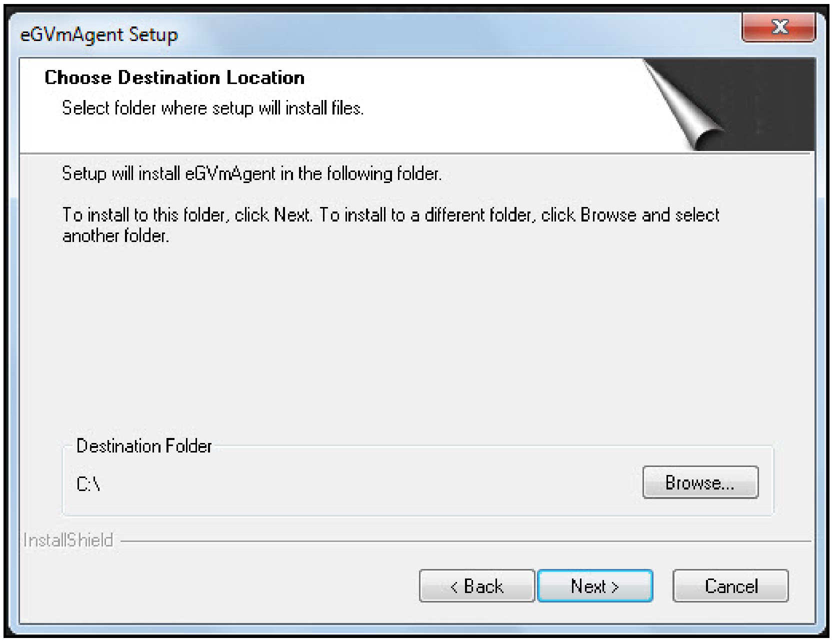

eGVmAgent Setup

Choose Destination Location

Select folder where setup will install files.

Setup will install eGVmAgent in the following folder.

To install to this folder, click Next. To install to a different folder, click Browse and select another folder.

Destination Folder

C:\

Figure 1.15: 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 1.16 to proceed.

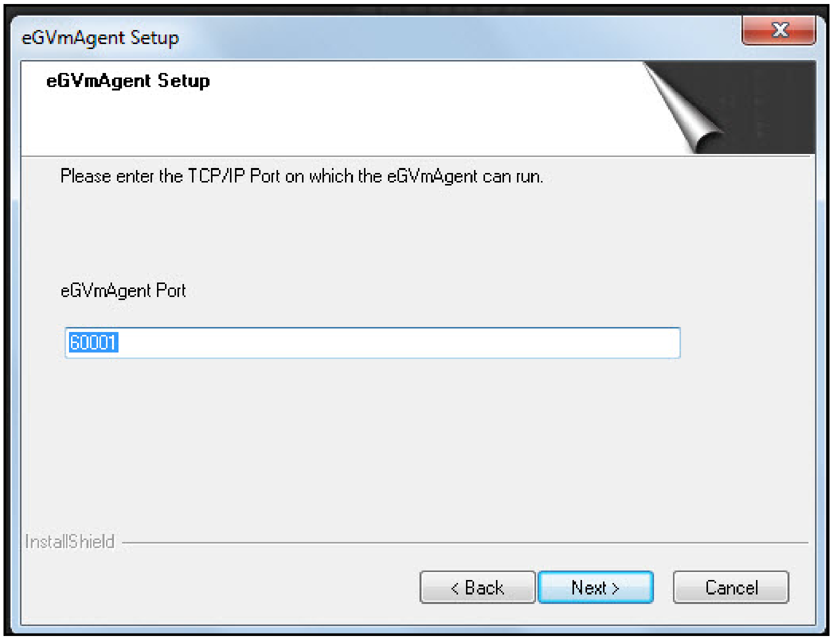

eGVmAgent Setup

eGVmAgent Setup

Please enter the TCP/IP Port on which the eGVmAgent can run.

eGVmAgent Port

60001

Figure 1.16: Specifying the VM agent port

6. A summary of your specifications then follows (see Figure 1.17). Click **Next** to proceed.

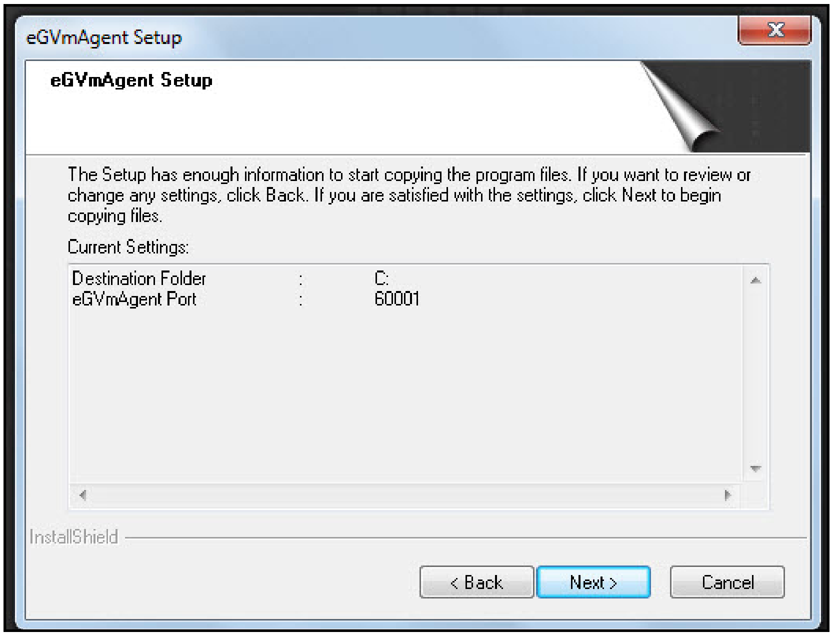

eGVmAgent Setup

eGVmAgent Setup

The Setup has enough information to start copying the program files. If you want to review or change any settings, click Back. If you are satisfied with the settings, click Next to begin copying files.

Current Settings:

```
Destination Folder      :      C:
eGVmAgent Port          :      60001
```

Figure 1.17: A summary of your specifications

7. Finally, click the **Finish** button in Figure 1.18 to complete the installation.

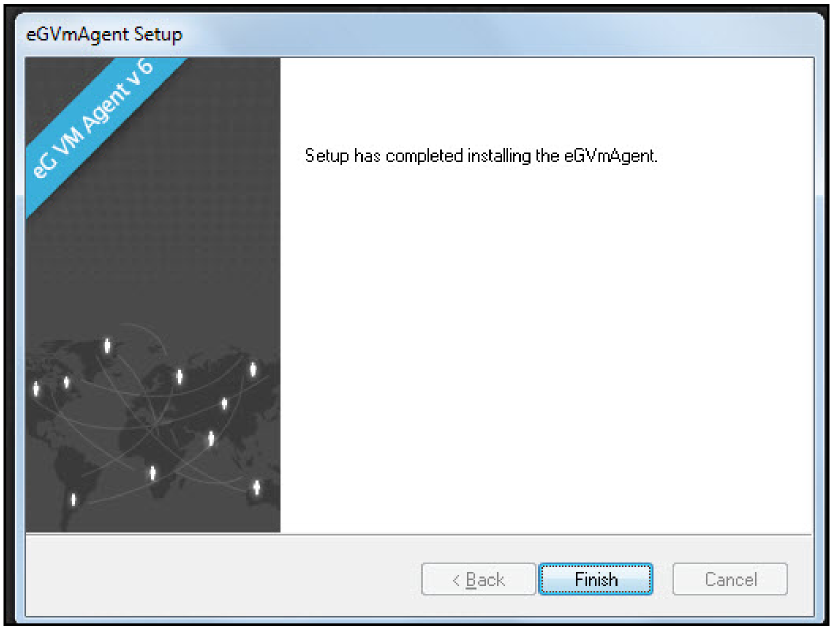

Setup has completed installing the eGVmAgent.

Figure 1.18: Finishing the installation

1.4.2 Silent Installation of the eG VM Agent

To silently install the eG VM agent on Windows VMs, follow the broad steps outlined below:

1. Creating silent mode script for eGVmagent installation
2. Installing eGVmAGent in silent mode

Each of these steps have been discussed elaborately below.

1.4.2.0.1 Creating a Silent Mode Script

For this, follow the procedure detailed below:

1. Login to a target Windows VM.
2. From the command prompt, run the following command to launch the normal mode installation of the eG VM Agent.

   *eGVMAgent_**<32/64>**.exe /a /r /f1**"<Full path to the script file into which the installation inputs will be stored>"***

   For example:

   eGVMAgent_x64.exe /a /r /f1"C:\script\eGVMAgent.iss"

3. Upon execution, this command will automatically create a script file of the given name in the location mentioned in the command.
4. Command execution will also begin the normal mode installation of the eG VM Agent. Provide inputs as

and when necessary to proceed with the installation.

5. These inputs will be automatically recorded in the script file that was created in step 3.

1.4.2.0.2 Installing the eG VM Agent in the Silent Mode

Follow the steps given below to install the eG VM Agent in the silent mode:

1. Login to the Windows VM where the script file containing the inputs for installation resides.
2. Copy the script file from this VM to the Windows VM on which you want to install the eG VM Agent in the silent mode.
3. Copy the eG VM Agent installation executable also to the target Windows VM.
4. Next, on the target Windows VM, run the following command from the command prompt:

   *eGVMAgent_**<32/64>**.exe /a /s /f1"**<Full path to the script file containing the inputs for the installation>**"*

   For example:

   eGVmAgent_x64.exe /a /s /f1"C:\script\eGVMAgent.iss"
5. Upon successful execution, this command will automatically install the eG VM Agent on the target Windows VM.
6. You can then repeat steps 1-5 on each Windows VM where you want to install the eG VM Agent.

1.4.3 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.

1.4.4 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.

1.4.5 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 VM 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.

1.5 Configuring the eG Agent to Monitor NVIDIA Graphics Processing Units (GPUs)

Citrix XenServer employs the NVIDIA GRID Virtual GPU (vGPU) technology to provide exceptional graphics performance for virtual desktops. NVIDIA GRID vGPU enables multiple Virtual Machines (VM) to have simultaneous, direct access to a single physical GPU, using the same NVIDIA graphics drivers that are deployed on non-virtualized Operating Systems. Under the control of NVIDIA's GRID Virtual GPU Manager, which runs in XenServer's Control Domain (dom0), GRID physical GPUs are capable of supporting multiple virtual GPU devices (vGPUs) that can be assigned directly to VMs.

To determine whether/not the vGPUs assigned to a VM are adequate for the graphic processing requirements of the that VM, administrators must understand whether/not memory-intensive I/O operations are performed on that VM and if so, how they impact vGPU usage. To perform this check, administrators can periodically run the **GPU – VM** test on every VM.

For this test to run and report metrics, the **NVWMI** must be installed on every VM.

NVIDIA WMI (NVWMI) is a graphics and display management and control technology that interfaces to Microsoft's Windows Management Instrumentation infrastructure, specific to NVIDIA graphics processing units (GPUs). This allows scripts and programs to be created that configure specific GPU related settings, perform automated tasks, retrieve and display a range of information related to the GPU as well as many other administrative tasks and functions.

The following NVIDIA products support NVWMI:

- NVIDIA Quadro® K600
- Quadro® K6000
- Quadro K5000
- Quadro K4000
- Quadro K2000D
- Quadro K2000
- Quadro FX 5800
- Quadro FX 580
- Quadro FX 570
- Quadro FX 5600
- Quadro FX 4800
- Quadro FX 4700 X2
- Quadro FX 4600
- Quadro FX 380 Low Profile
- Quadro FX 3800
- Quadro FX 380
- Quadro FX 3700
- Quadro FX 370
- Quadro FX 3450
- Quadro FX 1800
- Quadro FX 1700
- Quadro CX
- Quadro 7000
- Quadro 6000
- Quadro 600

- Quadro 5000
- Quadro 410
- Quadro 4000 for Mac
- Quadro 4000
- Quadro 400
- Quadro 2000D
- Quadro 2000
- NVIDIA® NVS® 510
- NVS 450
- NVS 420
- NVS 315
- NVS 310
- NVS 300
- NVS 295
- NVS 290
- Quadro Plex S Series
- Quadro Plex Model IV
- Quadro Plex D Series
- Quadro Plex 7000

NVWMI can be installed in the following three ways:

- When the NVIDIA GPU driver is installed
- Via a standalone install
- Via command line install

Note:

NVWMI is only supported on Windows 7 and later operating systems.

Each of these installation options are detailed in the sub-sections.

1.5.1 Installing NVWMI as Part of the NVIDIA GPU Driver

When installing the full GPU driver, please be aware that NVWMI is not installed by default. To ensure that it is installed, it is necessary to customize the installation and ensure that the NVWMI check box is checked, as shown in Figure 1.19 and Figure 1.20 in the following procedure.

1. Download the full GPU driver from **www.nvidia.com/drivers**.
2. Double click the **setup.exe**.

3. Select **Custom (Advanced Options)** under “Installation options".

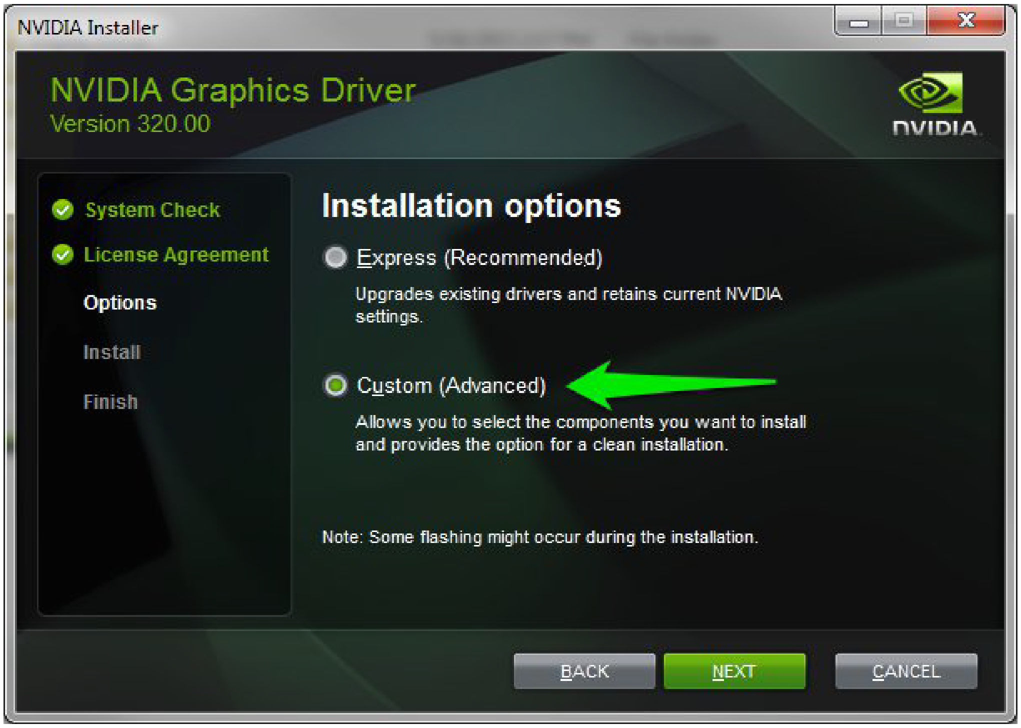

NVIDIA Installer

NVIDIA Graphics Driver
Version 320.00

Installation options

○ Express (Recommended)

Upgrades existing drivers and retains current NVIDIA settings.

● Custom (Advanced)

Allows you to select the components you want to install and provides the option for a clean installation.

Note: Some flashing might occur during the installation.

Figure 1.19: Custom check box

4. Select the *NVIDIA WMI* checkbox under **Custom installation options**.

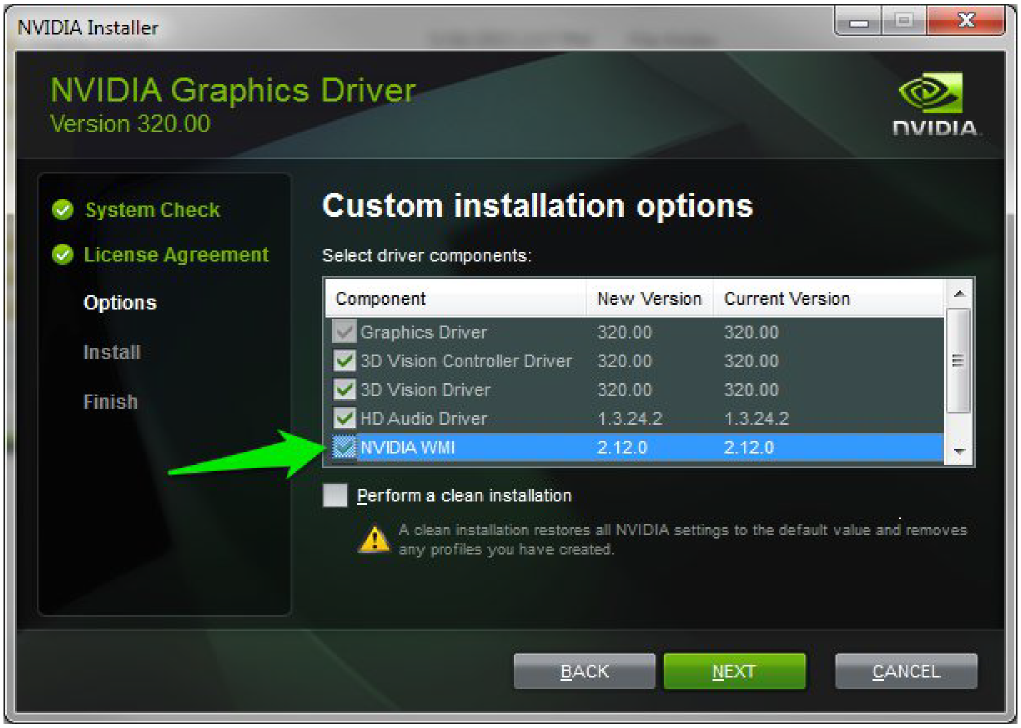

NVIDIA Installer

NVIDIA Graphics Driver
Version 320.00

- System Check
- License Agreement
- Options
- Install
- Finish

Custom installation options

Select driver components:

Component	New Version	Current Version
☑ Graphics Driver	320.00	320.00
☑ 3D Vision Controller Driver	320.00	320.00
☑ 3D Vision Driver	320.00	320.00
☑ HD Audio Driver	1.3.24.2	1.3.24.2
☑ NVIDIA WMI	2.12.0	2.12.0

□ Perform a clean installation

A clean installation restores all NVIDIA settings to the default value and removes any profiles you have created.

Figure 1.20: Selecting the NVIDIA WMI check box

5. Continue to install the driver and follow the on-screen instructions to complete the installation.
6. Repeat steps 1-5 on every VM on which you want to run the **GPU Stats – VM** test.

1.5.2 Installing NVWMI from the Standalone Installer

1. Go to **www.nvidia.com/drivers**.
2. Select the **NVIDIA Enterprise Manager** under **Beta, Older drivers and more**.

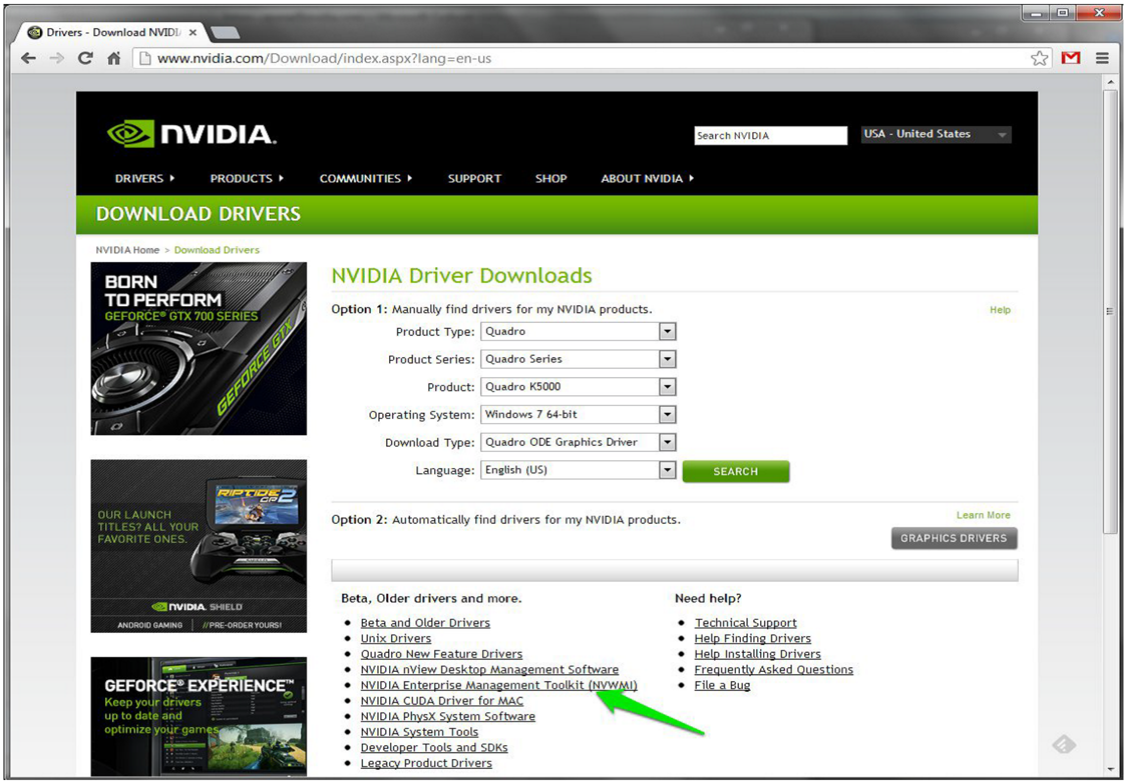
Drivers - Download NVIDI/
www.nvidia.com/Download/index.aspx?lang=en-us
NVIDIA
Search NVIDIA
USA - United States
DRIVERS
PRODUCTS
COMMUNITIES
SUPPORT
SHOP
ABOUT NVIDIA
DOWNLOAD DRIVERS
NVIDIA Home > Download Drivers
BORN TO PERFORM
GEFORCE® GTX 700 SERIES
OUR LAUNCH TITLES? ALL YOUR FAVORITE ONES.
NVIDIA SHIELD
ANDROID GAMING | //PRE-ORDER YOURS!
GEFORCE® EXPERIENCE™
Keep your drivers up to date and optimize your games

NVIDIA Driver Downloads

Option 1: Manually find drivers for my NVIDIA products. Help

Product Type: Quadro

Product Series: Quadro Series

Product: Quadro K5000

Operating System: Windows 7 64-bit

Download Type: Quadro ODE Graphics Driver

Language: English (US)

SEARCH

Option 2: Automatically find drivers for my NVIDIA products. Learn More

GRAPHICS DRIVERS

Beta, Older drivers and more.

- Beta and Older Drivers
- Unix Drivers
- Quadro New Feature Drivers
- NVIDIA nView Desktop Management Software
- NVIDIA Enterprise Management Toolkit (NVWMI)
- NVIDIA CUDA Driver for MAC
- NVIDIA PhysX System Software
- NVIDIA System Tools
- Developer Tools and SDKs
- Legacy Product Drivers

Need help?

- Technical Support
- Help Finding Drivers
- Help Installing Drivers
- Frequently Asked Questions
- File a Bug

Figure 1.21: NVIDIA Enterprise Management Toolkit (NVWMI)

3. Ensure to verify the minimum driver version is met and download the driver.

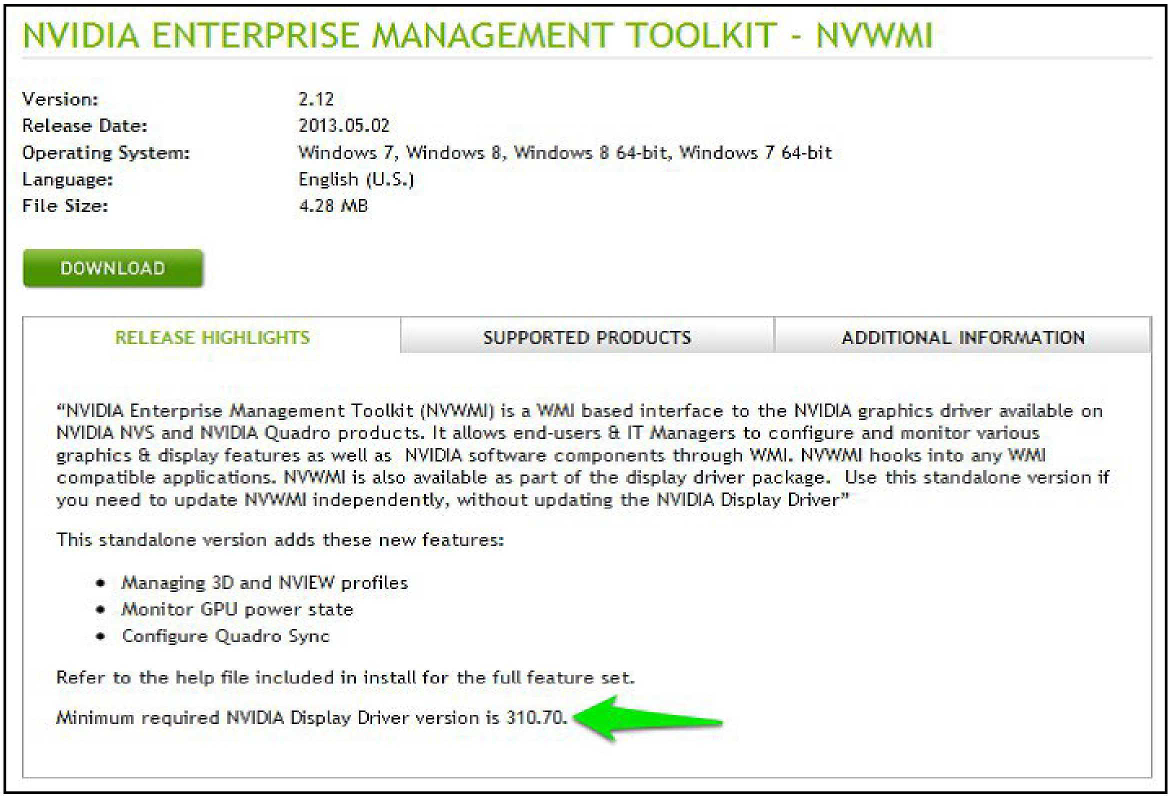

NVIDIA ENTERPRISE MANAGEMENT TOOLKIT - NVWMI

Version:	2.12
Release Date:	2013.05.02
Operating System:	Windows 7, Windows 8, Windows 8 64-bit, Windows 7 64-bit
Language:	English (U.S.)
File Size:	4.28 MB

"NVIDIA Enterprise Management Toolkit (NVWMI) is a WMI based interface to the NVIDIA graphics driver available on NVIDIA NVS and NVIDIA Quadro products. It allows end-users & IT Managers to configure and monitor various graphics & display features as well as NVIDIA software components through WMI. NVWMI hooks into any WMI compatible applications. NVWMI is also available as part of the display driver package. Use this standalone version if you need to update NVWMI independently, without updating the NVIDIA Display Driver"

This standalone version adds these new features:

- Managing 3D and NVIEW profiles
- Monitor GPU power state
- Configure Quadro Sync

Refer to the help file included in install for the full feature set.

Minimum required NVIDIA Display Driver version is 310.70.

Figure 1.22: NVIDIA WMI minimum driver requirements

4. Continue to install the driver and follow the on-screen instructions to complete the installation.
5. Repeat steps 1-4 on every VM on which you want to run the **GPU Stats – VM** test.

1.5.3 Installing NVWMI via the Command Line

1. Download the driver from the URL **www.nvidia.com/drivers**
2. Using the command line navigate to the extracted directory of the downloaded driver.
3. Issue the following command at the command prompt:

   setup.exe [switches][package]

   Here:

   - **[package]** = Display NVWMI
   - **[switches]** can be one of the following:

Switches	Description
-s	Silent install or uninstall

Switches	Description
-k	Force a reboot after install or uninstall
-uninstall	Performs an uninstall instead of an install

To confirm if NVIDIA WMI is already installed, simply look in the list of programs available to change or uninstall in the **Control Panel** under **Programs and Features** heading.

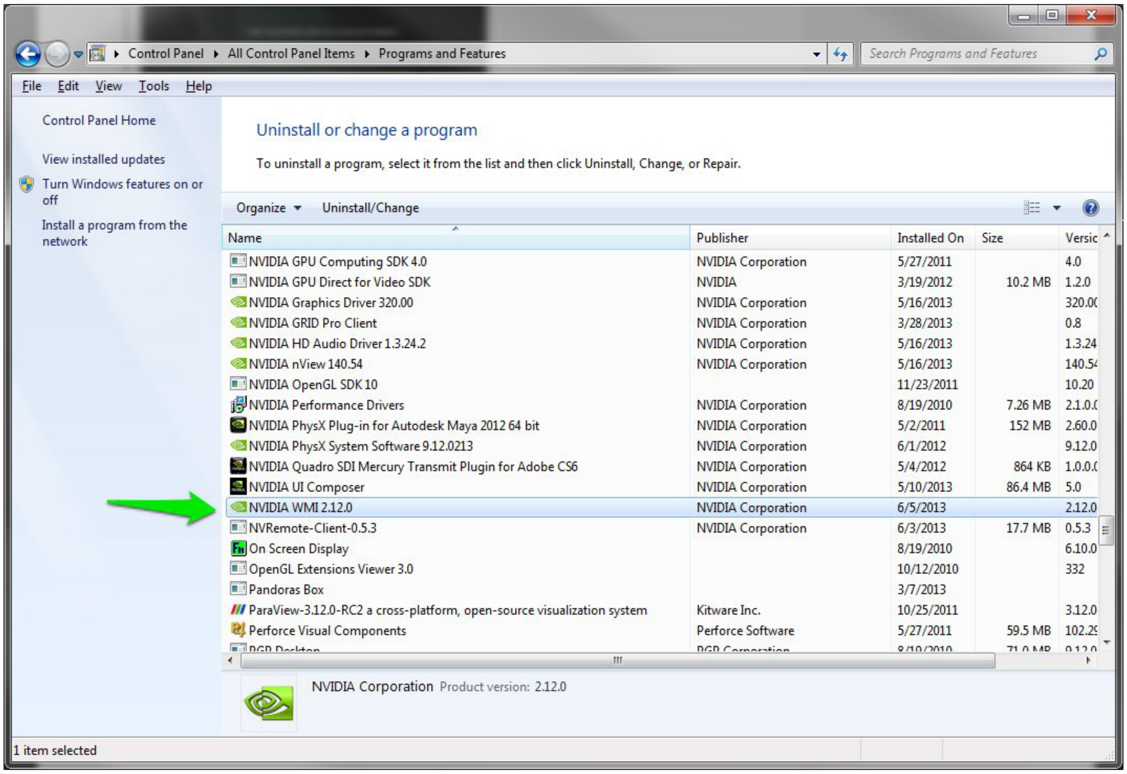
Control Panel ▸ All Control Panel Items ▸ Programs and Features
Search Programs and Features
File Edit View Tools Help
Control Panel Home
View installed updates
Turn Windows features on or off
Install a program from the network
Organize ▾
Uninstall/Change

Uninstall or change a program

To uninstall a program, select it from the list and then click Uninstall, Change, or Repair.

Name	Publisher	Installed On	Size	Versic
NVIDIA GPU Computing SDK 4.0	NVIDIA Corporation	5/27/2011		4.0
NVIDIA GPU Direct for Video SDK	NVIDIA	3/19/2012	10.2 MB	1.2.0
NVIDIA Graphics Driver 320.00	NVIDIA Corporation	5/16/2013		320.0
NVIDIA GRID Pro Client	NVIDIA Corporation	3/28/2013		0.8
NVIDIA HD Audio Driver 1.3.24.2	NVIDIA Corporation	5/16/2013		1.3.24
NVIDIA nView 140.54	NVIDIA Corporation	5/16/2013		140.5
NVIDIA OpenGL SDK 10		11/23/2011		10.20
NVIDIA Performance Drivers	NVIDIA Corporation	8/19/2010	7.26 MB	2.1.0.
NVIDIA PhysX Plug-in for Autodesk Maya 2012 64 bit	NVIDIA Corporation	5/2/2011	152 MB	2.60.0
NVIDIA PhysX System Software 9.12.0213	NVIDIA Corporation	6/1/2012		9.12.0
NVIDIA Quadro SDI Mercury Transmit Plugin for Adobe CS6	NVIDIA Corporation	5/4/2012	864 KB	1.0.0.
NVIDIA UI Composer	NVIDIA Corporation	5/10/2013	86.4 MB	5.0
NVIDIA WMI 2.12.0	NVIDIA Corporation	6/5/2013		2.12.0
NVRemote-Client-0.5.3	NVIDIA Corporation	6/3/2013	17.7 MB	0.5.3
On Screen Display		8/19/2010		6.10.0
OpenGL Extensions Viewer 3.0		10/12/2010		332
Pandoras Box		3/7/2013		
ParaView-3.12.0-RC2 a cross-platform, open-source visualization system	Kitware Inc.	10/25/2011		3.12.0
Perforce Visual Components	Perforce Software	5/27/2011	59.5 MB	102.2
RGP Desktop	RGP Corporation	8/19/2010	71.0 MB	9.12.0

NVIDIA Corporation Product version: 2.12.0

1 item selected

Figure 1.23: Checking the Control Panel for NVIDIA WMI

4. Repeat steps 1-3 on every VM on which you want to run the **GPU Stats – VM** test.

Once the pre-requisites for the “inside” and “outside” views are fulfilled, the eG agent periodically executes tests on the XenServer host and VMs, collects useful resource usage metrics from them, and reports the measurements so collected to the eG manager. The eG manager then consolidates and presents the metrics so received using a layer model representation.

The chapters that follow discuss the layer models that eG Enterprise offers out-of-the-box for the *Citrix XenServer* and *Citrix XenServer - VDI* components, and what each layer reveals.

2

Administering the eG Manager to monitor Citrix XenServer

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

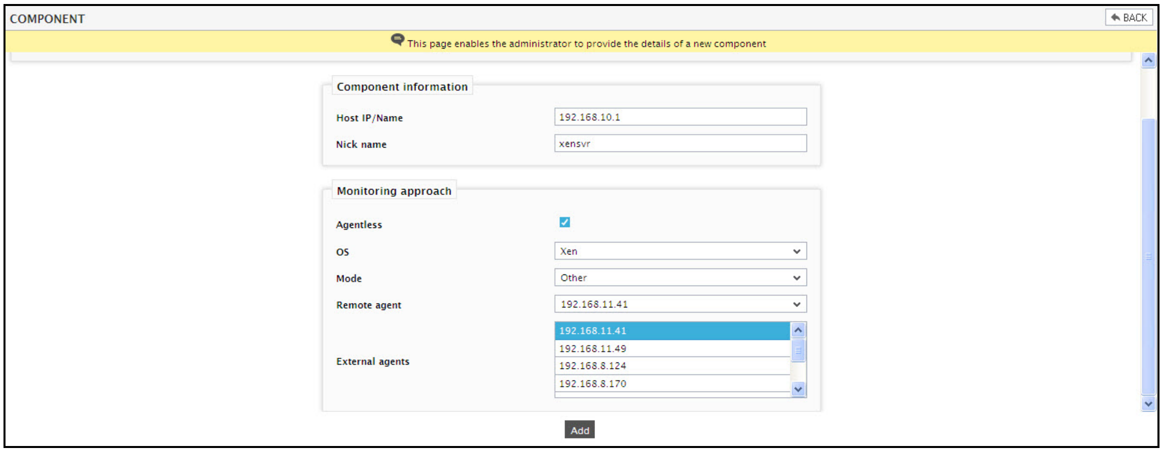

COMPONENT

This page enables the administrator to provide the details of a new component

Component information

Host IP/Name: 192.168.10.1

Nick name: xensvr

Monitoring approach

Agentless: ☑

OS: Xen

Mode: Other

Remote agent: 192.168.11.41

External agents:
192.168.11.41
192.168.11.49
192.168.8.124
192.168.8.170

Figure 2.1: Adding the Citrix XenServer

3. To monitor a Citrix XenServer, select **Xen** as the OS and set the Mode to **SSH**. Then, pick a **Remote agent** from Figure 2.1 to perform the agentless monitoring, and also assign an **External agent** to the server to verify its external availability and network health. Finally, click the **Add** button in Figure 2.1 to manage the component.
4. When you attempt to sign out, a list of unconfigured tests appears (see Figure 2.2).

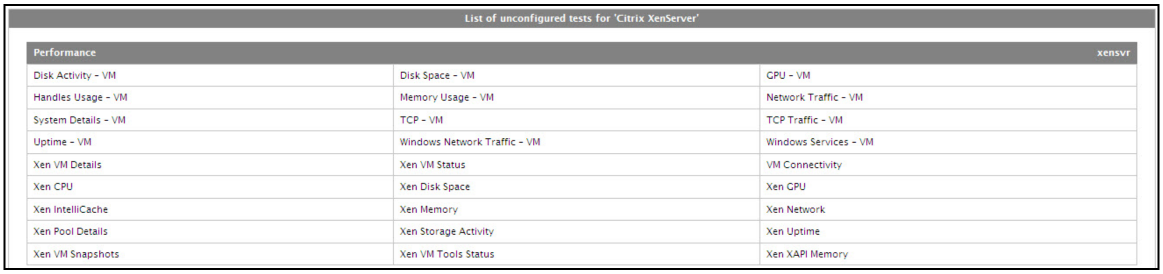

List of unconfigured tests for 'Citrix XenServer'

Performance		xensvr
Disk Activity - VM	Disk Space - VM	GPU - VM
Handles Usage - VM	Memory Usage - VM	Network Traffic - VM
System Details - VM	TCP - VM	TCP Traffic - VM
Uptime - VM	Windows Network Traffic - VM	Windows Services - VM
Xen VM Details	Xen VM Status	VM Connectivity
Xen CPU	Xen Disk Space	Xen GPU
Xen IntelliCache	Xen Memory	Xen Network
Xen Pool Details	Xen Storage Activity	Xen Uptime
Xen VM Snapshots	Xen VM Tools Status	Xen XAPI Memory

Figure 2.2: Figure 2.2: List of unconfigured tests for Citrix XenServer

5. Click on the **Disk Space - VM** test in Figure 2.2 to configure it. To know how to configure the test, click here.\
6. Once all the tests are configured, signout of the eG administrative interface

2

Administering the eG Manager to Monitor Citrix XenServer - VDI

1. Log into the eG administrative interface.
2. eG Enterprise cannot automatically discover the Citrix XenServer - VDI. You need to manually add the Citrix XenServer - VDI using the **COMPONENTS** page (see Figure 2.3) that appears when the Infrastructure -> Components -> Add/Modify menu sequence is followed. Remember that components manually added are managed automatically.

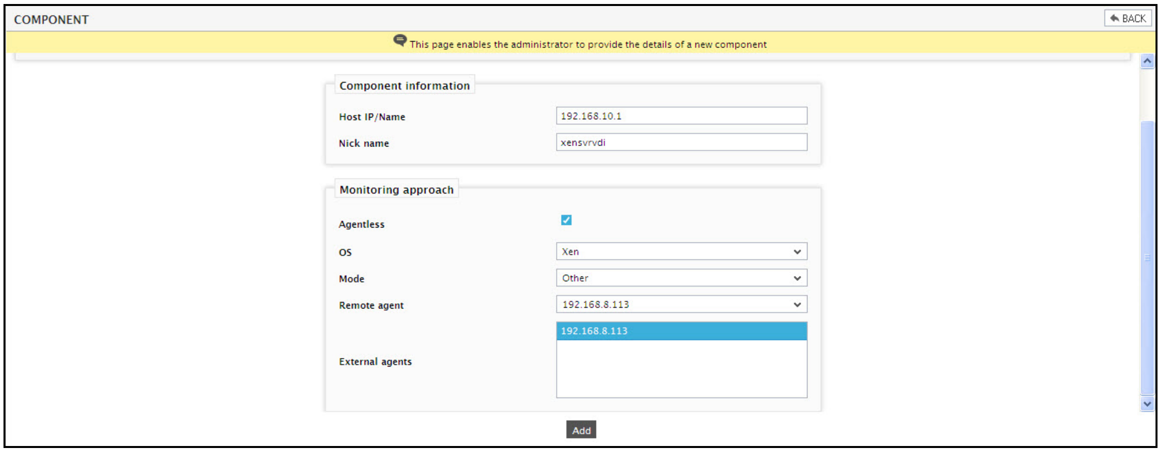

COMPONENT

This page enables the administrator to provide the details of a new component

Component information

Host IP/Name: 192.168.10.1

Nick name: xensvrvdi

Monitoring approach

Agentless: ☑

OS: Xen

Mode: Other

Remote agent: 192.168.8.113

External agents: 192.168.8.113

Figure 2.3: Adding the Citrix XenServer – VDI

3. To monitor a Citrix XenServer - VDI, select **Xen** as the **OS** and set the **Mode** to **SSH**. Then, pick a **Remote agent** from Figure 2.3 to perform the agentless monitoring, and also assign an **External agent** to the server to verify its external availability and network health. Finally, click the **Add** button in Figure 2.3 to manage the component.
4. When you attempt to sign out, a list of unconfigured tests appears (see Figure 2.4).

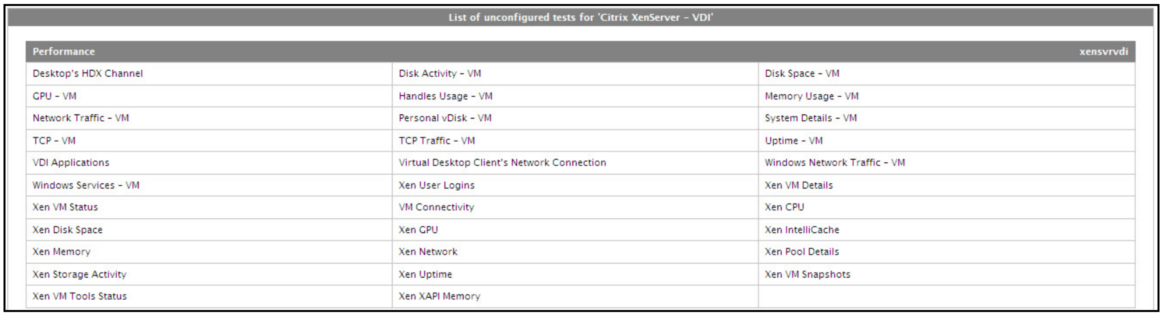

List of unconfigured tests for 'Citrix XenServer - VDI'

Performance		xensvrvdi
Desktop's HDX Channel	Disk Activity - VM	Disk Space - VM
GPU - VM	Handles Usage - VM	Memory Usage - VM
Network Traffic - VM	Personal vDisk - VM	System Details - VM
TCP - VM	TCP Traffic - VM	Uptime - VM
VDI Applications	Virtual Desktop Client's Network Connection	Windows Network Traffic - VM
Windows Services - VM	Xen User Logins	Xen VM Details
Xen VM Status	VM Connectivity	Xen CPU
Xen Disk Space	Xen GPU	Xen IntelliCache
Xen Memory	Xen Network	Xen Pool Details
Xen Storage Activity	Xen Uptime	Xen VM Snapshots
Xen VM Tools Status	Xen XAPI Memory	

Figure 2.4: List of unconfigured tests for Citrix XenServer - VDI

5. Click on the **Desktop's HDX Channel** test in Figure 2.4 to configure it. To know how to configure the test, click here.

2

The Citrix XenServer Monitoring Model

As stated earlier, you can use the *Citrix XenServer* monitoring model for monitoring the performance of XenServers hosting one/more server applications. Figure 2.5 depicts the *Citrix XenServer* monitoring model.

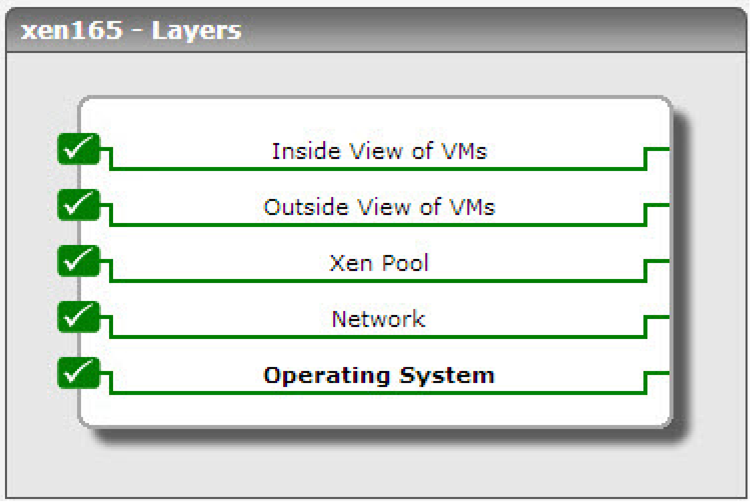

Figure 2.5: The Citrix XenServer Monitoring Model

Each layer depicted by Figure 2.5 above reports a wide variety of statistics, which enable administrators to find quick and accurate answers for the following questions:

- How many guest VMs are running on each XenServer, what is the IP address of each of guests, what operating system is each guest running, and when was the guest powered on?
- How much memory is allocated to each guest and does each guest VM have sufficient free memory?

- Does the XenServer have sufficient memory available to support the guest VMs that it is hosting?
- What is the CPU utilization of the XenServer and which of the guest VMs is taking up excessive CPU?
- Which application(s) running on each of the guest VMs is taking CPU, memory, and disk resources?
- Is there sufficient disk space in each of the disk partitions of the guest operating system?
- Which of the guests is seeing the highest and lowest network traffic?
- Is there excessive queuing for disk access on any of the guest VMs?

The sections that follow will discuss each of these layers in great detail.

2.1 The Operating System Layer

The tests mapped to this layer report the percentage of physical resources used by the control domain and the XenServer host.

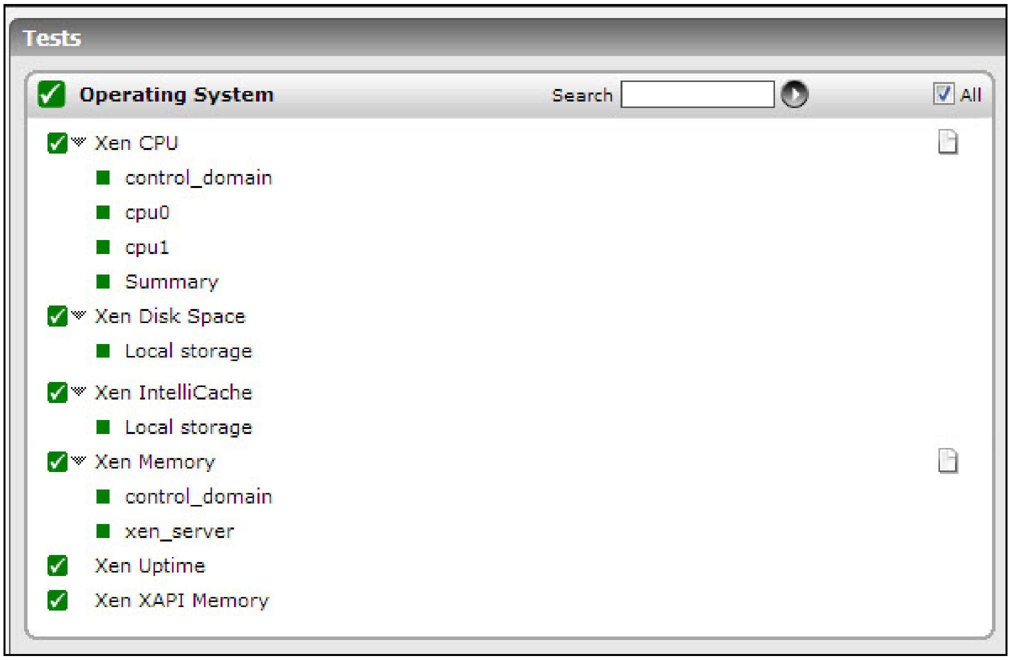

Figure 2.6: The tests mapped to the Operating System layer

2.1.1 CPU - Xen Test

This test reports the real-time CPU utilization statistics pertaining to the control domain and every processor supported by the XenServer. The control domain is a privileged VM that provides low-level services to other VMs, such as providing access to physical devices. It also runs the management tool stack. Using this test, administrators can determine whether there is a resource-contention at the XenServer host, and if so, where - at the control domain? or with the processors supported by the host?

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each processor supported by the XenServer host and for the control domain executing on the host

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. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the **XEN USER** text box. Root user privileges are mandatory when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the root user, then, you have the option of configuring a user with *pool-admin* privileges as the **XEN USER**. If you do not want to expose the credentials of a root/pool-admin user, then you can configure the tests with the credentials of a **XEN USER** with *Read-only* privileges to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the control domain descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a xen user who has *Read-only* privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via **SSH**.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
4. **XEN PASSWORD** - The password of the specified xen user needs to be mentioned here.
5. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
6. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

7. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the webport parameter is set to 80 or 443 depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

8. **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 abnormalfrequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Physical CPU usage:	Indicates the percentage of physical CPU used by this processor or the control domain.	Percent	A very high value for this measure indicates excessive CPU utilization by that processor or control domain. 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. If adequate CPU resources are not available to the host, then the guests too would experience a resource crunch; this in turn, could cause significant slowdowns in the execution of applications hosted by the guests. A high value for the control_domain descriptor indicates that one/more control domain processes are consuming CPU resources

Measurement	Description	Measurement Unit	Interpretation
			excessively.

The detailed diagnosis capability that is available only for the *control_domain* descriptor, will list the CPU-intensive processes executing on the control domain. Resource intensive processes can be killed/stopped to avoid excessive CPU usage. Alternatively, administrators can allocate additional CPU resources to the host to ensure that processes have adequate memory for execution.

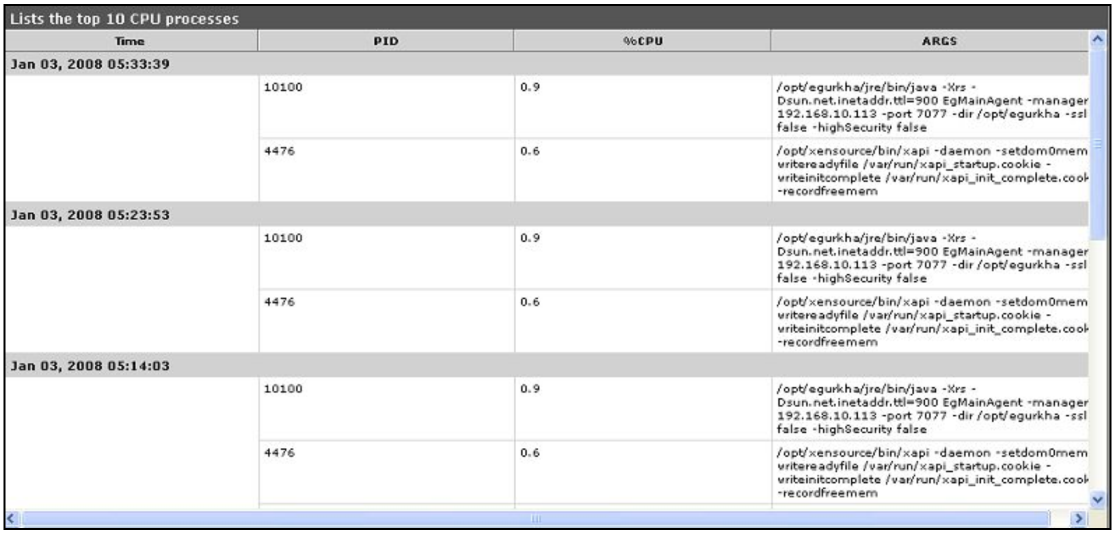

Lists the top 10 CPU processes

Time	PID	%CPU	ARGS
Jan 03, 2008 05:33:39			
	10100	0.9	/opt/egurkha/jre/bin/java -Xrs -Dsun.net.inetaddr.ttl=900 EgMainAgent -manager 192.168.10.113 -port 7077 -dir /opt/egurkha -ssl false -highSecurity false
	4476	0.6	/opt/xensource/bin/xapi -daemon -setdom0mem writereadyfile /var/run/xapi_startup.cookie -writeinitcomplete /var/run/xapi_init_complete.cook -recordfreemem
Jan 03, 2008 05:23:53			
	10100	0.9	/opt/egurkha/jre/bin/java -Xrs -Dsun.net.inetaddr.ttl=900 EgMainAgent -manager 192.168.10.113 -port 7077 -dir /opt/egurkha -ssl false -highSecurity false
	4476	0.6	/opt/xensource/bin/xapi -daemon -setdom0mem writereadyfile /var/run/xapi_startup.cookie -writeinitcomplete /var/run/xapi_init_complete.cook -recordfreemem
Jan 03, 2008 05:14:03			
	10100	0.9	/opt/egurkha/jre/bin/java -Xrs -Dsun.net.inetaddr.ttl=900 EgMainAgent -manager 192.168.10.113 -port 7077 -dir /opt/egurkha -ssl false -highSecurity false
	4476	0.6	/opt/xensource/bin/xapi -daemon -setdom0mem writereadyfile /var/run/xapi_startup.cookie -writeinitcomplete /var/run/xapi_init_complete.cook -recordfreemem

Figure 2.7: The detailed diagnosis of the Physical CPU usage measure

Note:

In case of multi-processor systems, this test typically reports the CPU usage for every processor. The **Summary** descriptor reports the average CPU usage across processors.

2.1.2 Disk Space – Xen Test

This test reports the space usage of every volume group supported by the XenServer host.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each volume group on a XenServer host

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. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of

interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the **XEN USER** text box. Root user privileges are mandatory when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the root user, then, you have the option of configuring a user with *pool-admin* privileges as the **XEN USER**. If you do not want to expose the credentials of a root/pool-admin user, then you can configure the tests with the credentials of a **XEN USER** with *Read-only* privileges to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the control domain descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a xen user who has *Read-only* privileges to the XenServer:

- Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
- In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
- Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via **SSH**.
- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

4. **XEN PASSWORD** - The password of the specified xen user needs to be mentioned here.
5. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
6. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.
7. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the webport parameter is set to 80 or 443 depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the

exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total physical extents:	Indicates the total number of physical extents on this volume group	Number	For the **Total** descriptor, this measure will report the total number of physical extents across all volume groups on the XenServer host.
Used physical extents:	Indicates the number of physical extents allocated to this volume group.	Number	For the **Total** descriptor, this measure will report the total number of physical extents allocated to all volume groups on the XenServer host.
Free physical extents:	Indicates the number of allocated extents that are still unused by this volume group.	Number	For the **Total** descriptor, this measure will report the total number of unused physical extents in all volume groups on the XenServer host. Ideally, the value of this measure should be high.
Physical extent size:	Indicates the current size of the physical extents in this volume group.	MB	For the **Total** descriptor, this measure will report the total size of the physical extents in all volume groups on the XenServer host.
Disk space in use by VMs:	Indicates the space in this volume group that is used by VMs.	GB	For the **Total** descriptor, this measure will report the total space used across all volume groups on the XenServer host. A high value indicates excessive utilization of a volume group. If the value of this measure is high for the **Total** descriptor, it indicates that the usage across volume groups is very high.
Space available for	Indicates the amount of free	GB	This measure reports the difference

Measurement	Description	Measurement Unit	Interpretation
allocation to new VMs:	space on this volume group.		between the *Total disk capacity* measure and the Disk space in use by VMs measure. Ideally, the value of this measure should be high. A very low value could indicate that one/more vMs are utilizing the allocated disk space excessively. To confirm this, check the value of the *Disk space in use by VM*s measure; if this value is very high, it indicates that the VMs are eroding disk space. For the **Total** descriptor, this measure will report the total disk space that is free across all volume groups, and can be allocated to new VMs.
Usage of allocated space:	Indicates the percentage of space allocated from this volume group that is used by VMs.	Percent	A value close to 100% is a cause for concern, as it indicates that one/more VMs are draining the disk space. If the situation persists or aggravates, VMs will experience a severe space crunch, which can bring all VM operations to a halt. To avoid this, you will have to isolate the VMs that are eroding the allocated disk space, clear unnecessary files from the disks of those VMs, and make additional space. Else, you should allocate more space to such VMs. For the **Total** descriptor, this measure will report the percentage of the total space allocations from all volume groups that is currently utilized by all VMs on the host. If the **Total** descriptor of this measure reports a high value, it indicates that the VMs on the host are consuming too much space that is allocated to them. You can then compare the value of this measure

Measurement	Description	Measurement Unit	Interpretation
			across volume groups to know which volume group is being used excessively.
Usage of disk capacity:	Indicates the percentage of the capacity of this volume group that is currently in use.	Percent	Ideally, the value of this measure should be low. A very high value indicates that the volume group is running out of space. Under such circumstances, you may want to check the value of Disk capacity allocated to VMs and the **Disk space in use by VMs** measures to know what is causing the space erosion – is it due to over-allocation of space to VMs? Or abnormal usage of space by the VMs? For the **Total** descriptor, this measure will report the percentage of the total disk capacity across all volume groups that is currently used. If the **Total** descriptor of this measure reports a high value, it indicates excessive usage of storage space across all volume groups. You can then compare the value of this measure across volume groups to know which volume group is being used excessively.
Total disk capacity:	Indicates the total capacity of this volume group.	GB	For the **Total** descriptor, this measure will report the total disk capacity of the host across all volume groups.
Disk capacity allocated to VMs:	Indicates the total disk space in this volume group that has been allocated to all VMs on the host.	GB	For the **Total** descriptor, this measure will report the total space that is allocated to all VMs from all volume groups on the host.
Disk capacity available for allocation:	Indicates the total disk space in this volume group that is yet to be allocated to	GB	For the **Total** descriptor, this measure will report the sum of the space in all volume groups that is still to be

Measurement	Description	Measurement Unit	Interpretation
	VMs.		allocated to VMs.
Disk space allocated to VMs:	Indicates the percentage of disk space in this volume group that has been allocated to VMs.	Percent	A very high value for this measure indicates that too much disk space has been allocated to the VMs. In such a situation, the host may not have enough disk space for its own operations, and this may cause the host's performance to deteriorate. For the **Total** descriptor, this measure will report the percentage of the total disk space across all volume groups that is allocated to all VMs on the host. If the **Total** descriptor of this measure reports a high value, it indicates that overall, too much storage space has been allocated to VMs. You can then compare the value of this measure across volume groups to know from which volume group the maximum space has been allocated.
Usage of allocated space:	Indicates the percentage of space allocated from this volume group that has been used by the VMs.	Percent	A high value for this measure indicates excessive usage of a volume group by the VMs. For the **Total** descriptor, this measure will indicate what percentage of the total allocated disk space (across all volume groups) that has been utilized by the VMs.

2.1.3 Xen Storage Activity Test

XenServer provides support for a broad range of storage hardware. The term Storage Repository (SR) is used to describe a particular storage target on which Virtual Disk Images (VDIs) are stored. A VDI is a disk abstraction that contains the contents of a disk as presented to a virtual machine. XenServer allows these VDIs to be supported on a large number of SR types, including local disks, NFS filers, Fibre Channel disks and shared iSCSI LUNs. The SR abstraction allows advanced storage features such as thin provisioning, VDI snapshots, and fast cloning to be exposed on storage targets that support them.

If a XenServer host is unable to or takes too much time to read from or write to an SR, it can result in undue delays in the provisioning and maintenance (i.e., creation, deletion, cloning, connecting, resizing, etc.) of virtual disk images. This, in turn, can significantly slowdown VM accesses. To ensure that the user experience with VMs remains top-notch, administrators should continuously monitor the I/O throughput of each storage repository (SR) supported by a XenServer host and quickly isolate the slow SRs. This is where the **Xen Storage Activity** test helps. By continuously measuring and reporting how well each SR handles read and write requests, this test precisely pinpoints slow SRs, thus prompting administrators to probe into the reasons for the slowness and fix them.

Note:

The performance metrics reported by this test are enabled by default in the **XenServer 6.1.0 Performance and Monitoring Supplemental Pack**. In XenServer 6.2.0 however, these metrics, though part of the core product, are disabled by default, owing to performance reasons related to XenCenter. This means that, when monitoring XenServer 6.2.0, this test will not report any metrics by default. In such cases, to make sure that the test reports metrics, do the following:

- Login to the XenServer host as **root** user.
- Enable the metrics by issuing the following command from the CLI:

  xe-enable-all-plugin-metrics true

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each PDB connecting the monitored XenServer host to an SR

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. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the **XEN USER** text box. Root user privileges are mandatory when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the root user, then, you have the option of configuring a user with *pool-admin* privileges as the **XEN USER**. If you do not want to expose the credentials of a root/pool-admin user, then you can configure the tests with the credentials of a **XEN USER** with *Read-only* privileges to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the control domain descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a xen user who has *Read-only* privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.

- In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
- Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via **SSH**.
- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

4. **XEN PASSWORD** - The password of the specified xen user needs to be mentioned here.
5. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
6. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.
7. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the webport parameter is set to 80 or 443 depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total throughput:	Indicates the throughput of this SR.	MB/Sec	A high value indicates high throughput and rapid I/O processing by the SR. Compare the value of this measure across SRs to identify the SR with the lowest throughput.

Measurement	Description	Measurement Unit	Interpretation
Read rate:	Indicates the rate at which the host reads data from this SR.	MB/Sec	Ideally, the value of this measure should be high. A consistent drop in the value of this measure indicates a reading bottleneck in the SR. You can compare the value of this measure across SRs to identify that SR which is the slowest in processing read requests.
Write rate:	Indicates the rate at which the host writes data to this SR.	MB/Sec	Ideally, the value of this measure should be high. A consistent drop in the value of this measure indicates a writing bottleneck in the SR. You can compare the value of this measure across SRs to identify that SR which is the slowest in processing write requests.
Total IOPS:	Indicates the rate at which I/O operations are performed by this SR.	Requests/Sec	This measure is a good indicator of the I/O processing capacity of the SR. A high value is hence desired for this measure. A consistent drop in this value could indicate a processing bottleneck. In such a situation, you can compare the value of the Read operations and Write operations measures of the corresponding SR to figure out where the bottleneck lies – in reading data from the SR? or in writing to the SR?
Read operations:	Indicates the rate at which this SR services read requests.	Requests/Sec	Ideally, the value of this measure should be high. A steady drop in this value indicates a slowdown in processing read requests. Compare the value of this measure across SRs to know which SR is the slowest in responding to read requests.
Write operations:	Indicates the rate at which this SR services write requests.	Requests/Sec	Ideally, the value of this measure should be high. A steady drop in this value indicates a slowdown in processing write requests. Compare the value of

Measurement	Description	Measurement Unit	Interpretation
			this measure across SRs to know which SR is the slowest in responding to write requests.
Time spent waiting for I/O:	Indicates the percentage of time the host's CPU was waiting for this SR to complete I/O processing.	Percent	A high value for this measure indicates that the SR is taking too long to complete I/O processing. This hints at a probable processing bottleneck with the SR.
Average latency:	Indicates the average time taken by this SR to process I/O requests.	MSecs	A high value for this measure is a cause for concern, as it indicates that the SR is highly latent and takes too long to process I/O. Compare the value of this measure across SRs to identify the most latent SR.
Average queue size:	Indicates the average number of I/O requests to this SR that are in queue for processing.	Number	If the value of this measure grows consistently, it indicates that the SR is unable to process requests quickly enough to clear the queue. The SR with the maximum number of queued requests could be experiencing a serious I/O processing bottleneck. To identify this SR, compare the value of this measure across SRs.
Current requests in flight:	Indicates the number of I/O requests to this SR that are currently being processed.	Number	

2.1.4 Xen IntelliCache Test

IntelliCache is a XenServer feature that can be used in a XenDesktop deployment to cache temporary and non-persistent operating-system data on the local XenServer host. IntelliCache is available for Machine Creation Services (MCS)-based desktop workloads that use NFS storage.

In a typical XenDesktop configuration (without IntelliCache), desktop VMs read the operating-system data from a master image on a costly shared storage array. When IntelliCache is enabled, a portion of the virtualmachine runtime reads and writes occur on low-cost local storage: XenServer caches the operating-system files on its local hard drive in a Read Cache.

Likewise, when IntelliCache is enabled, each desktop VM writes to its own Write Cache on the local host, preventing writes to shared storage. The read and write caches, if adequately sized and effectively used, can thus go a long way in reducing the load on the remote storage and the amount of network traffic. But, how does an administrator determine whether these caches are sized and used right? This can be achieved using the **Xen IntelliCache** test. This test continuously tracks the usage of the read and write caches in each of the local storage repositories and reports whether/not the caches are able to service most of the I/O requests to the shared storage, without actually accessing the shared storage! This way, administrators will be able to figure out whether/not the caches effectively serve the purpose they were intended for. In addition, the test also checks and reports the current cache size, so that administrators can judge its adequacy and make changes, if required.

Note:

The performance metrics reported by this test are enabled by default in the **XenServer 6.1.0 Performance and Monitoring Supplemental Pack**. In XenServer 6.2.0 however, these metrics, though part of the core product, are disabled by default, owing to performance reasons related to XenCenter. This means that, when monitoring XenServer 6.2.0, this test will not report any metrics by default. In such cases, to make sure that the test reports metrics, do the following:

- Login to the XenServer host as **root** user.
- Enable the metrics by issuing the following command from the CLI:

  xe-enable-all-plugin-metrics true

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each local SR on the monitored XenServer host

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. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the **XEN USER** text box. Root user privileges are mandatory when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the root user, then, you have the option of configuring a user with *pool-admin* privileges as the **XEN USER**. If you do not want to expose the credentials of a root/pool-admin user, then you can configure the tests with the credentials of a **XEN USER** with *Read-only* privileges to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the control domain descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a xen user who has *Read-only* privileges to the XenServer:
   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the

Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.

- In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
- Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via **SSH**.
- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

4. **XEN PASSWORD** - The password of the specified xen user needs to be mentioned here.
5. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
6. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.
7. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the webport parameter is set to 80 or 443 depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
IntelliCache hits:	Indicates the rate at which the caches on this SR serviced I/O requests.	Hits/Sec	A high value is desired for this measure. A low value is indicative of ineffective cache usage, which is typically caused

Measurement	Description	Measurement Unit	Interpretation
			if the cache is unable to serve most I/O requests. This in turn increases the load on the shared storage.
IntelliCache misses:	Indicates the number of requests per second that the caches on this SR did not service.	Misses/Sec	A low value is desired for this measure. A high value is indicative of ineffective cache usage, which is typically caused if the cache is unable to serve most I/O requests. This in turn increases the load on the shared storage.
IntelliCache hit ratio:	Indicates the percentage of I/O requests that were serviced by the caches on this SR.	Percent	A value over 80% for this measure is a sign of a healthy cache. A very low value indicates that the cache was unable to service most of the I/O requests to the SR. One of the common reasons for this is the lack of adequate space in the SR to store cached data. If the caches are not adequately sized, the read cache will not be able to hold frequently-referenced data and the write caches will not be able to accept data written to it. Check the value of the IntelliCache size measure to know how much space has been allocated to the caches on an SR. Should the local storage reach capacity, IntelliCache will transparently “fall back” to shared storage, increasing the processing overheads in the bargain. Forecasting your local disk space requirements helps prevent XenServer from having to fall back to shared storage to handle the IOPS demand. Depending on variables in your environment, like patterns of user activity, you may need to plan more space for your Read Cache size. In addition, your disk - space

Measurement	Description	Measurement Unit	Interpretation
			requirements could increase any time multiple catalogs are present, such as during an upgrade rollout. For example, if virtual machines use multiple versions of the same catalog, Read Cache space usage will increase proportionately. From a planning perspective, you should assume all of the master image could potentially be stored in the Read Cache. Consequently, if you have multiple catalogs on a host, you should assume that each catalog's master image could be stored in the Read Cache. For example, if you have two catalogs each with different versions of applications in them, both master images could potentially be stored. Likewise, if you are rolling out an operating system update, you may have two catalogs before users reboot and switch over to the new image.
IntelliCache size:	Indicates the size of the caches on this SR.	MB	

2.1.5 Xen XAPI Memory Test

XAPI – short for Xen API - is at the heart of the Citrix XenServer. It manages everything - all the resources in your XenServer environment. It reads in configs, initializes networking paths, checks licensing, maintains a database with all this info (the XAPI database), manages Storage Repositories, manages VMs, tracks VM states, manages networking, High Availability and other XenServer hosts in the pool and more. In short, every XenServer request or action is done via XAPI.

For performing all these operations, the XAPI daemon uses the memory allocated. If users to the XenServer or its VMs complain of slowdowns, it would be wise to check the memory usage of the XAPI daemon as well, amongst other things, as significant memory usage by the XAPI can also cause the host and VMs to choke. To enable administrators to track the memory usage of the XAPI daemon, the eG agent periodically runs the **Xen XAPI Memory** test. This test checks the memory usage of the XAPI daemon at configured intervals, captures abnormal usage patterns (if any), and proactively alerts administrators to them, so that potential slowdowns in the performance of the XenServer can be averted.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each local SR on the monitored XenServer host

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. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the **XEN USER** text box. Root user privileges are mandatory when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the root user, then, you have the option of configuring a user with *pool-admin* privileges as the **XEN USER**. If you do not want to expose the credentials of a root/pool-admin user, then you can configure the tests with the credentials of a **XEN USER** with *Read-only* privileges to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the control domain descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a xen user who has *Read-only* privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via **SSH**.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
4. **XEN PASSWORD** - The password of the specified xen user needs to be mentioned here.
5. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
6. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use

by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

- Obtain the server-certificate for the XenServer
- Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

7. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to 80 or 443 depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Memory usage:	Indicates the amount of allocated memory that is used by the XAPI daemon.	MB	A low value is desired for this measure.
Free memory:	Indicates the amount of allocated memory that is still unused by the XAPI daemon.	MB	Ideally, the value of this measure should be high. A consistent decrease in this value could indicate excessive memory usage by the XAPI.
Live memory:	Indicates the amount of live memory that is used by the XAPI.	MB	
Memory allocation:	Indicates the amount of memory allocated by the XAPI daemon.	MB	

2.1.6 Memory - Xen Test

When computing the memory footprint of a Xen host, there are two components that must be taken into consideration. First, there is the memory consumed by the Xen hypervisor itself. The hypervisor is a software program that manages multiple operating systems (or multiple instances of the same operating system) on a single computer system. It manages the system's processor, memory, and other resources to allocate what each operating system requires. Hypervisors are designed for a particular processor architecture and may also be called virtualization managers.

The other component that can impact the memory consumption of a XenServer host is the control domain. The control domain is a privileged VM that provides low-level services to other VMs, such as providing access to physical devices. It also runs the management tool stack.

The **Memory – Xen** test reports the memory usage of the control domain and hypervisor of the XenServer, and enables administrators to judge whether/not adequate free memory is available on the XenServer host for use by the virtual machines.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results each for the **control_domain** and the **xen_server**

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is *NULL*.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the **XEN USER** text box. Root user privileges are mandatory when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the root user, then, you have the option of configuring a user with *pool-admin* privileges as the **XEN USER**. If you do not want to expose the credentials of a root/pool-admin user, then you can configure the tests with the credentials of a **XEN USER** with *Read-only* privileges to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the control domain descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a xen user who has *Read-only* privileges to the XenServer:

    - Modify the target XenServer’s configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
    - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
    - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to

connect to the XenServer console via **SSH**.

- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified xen user needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to 80 or 443 depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.
9. **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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Free memory:	Indicates the amount of memory remaining unused with the host.	MB	Ideally, the value of this measure should be very high. If the value of this measure is very low or is steadily decreasing, it is indicative of a severe memory erosion at the host-level. If the free memory is very low or is decreasing consistently, you may want to compare the value of the *Memory in use* measure between the **xen_ server** and **control_ domain** descriptors of this test, to know where the host's memory has been spent the maximum - in supporting the control domain's processes? or because of memory-intensive VMs? The detailed diagnosis of this measure for the **control_ domain** descriptor reveals the top-10 memory consuming processes on the control domain and the amount and percentage of memory they consume. The detailed diagnosis of this measure for the **xen_server** descriptor reveals the top-10 memory consuming VMs and the amount and percentage of memory they consume.
Total memory available:	Indicates the total physical memory available for use at the host.	MB	
Memory allocated to VMs:	Indicates the amount of physical memory that is allocated to the VMs.	MB	This measure is only available for the xen_server descriptor of this test.
Available memory allocated to VMs:	Indicates the percentage of physical memory that is allocated to VMs.	Percent	This measure is only available for the xen_server descriptor of this test. A high value could indicate excessive memory allocation to the VMs. You can use the detailed diagnosis of

Measurement	Description	Measurement Unit	Interpretation
			this measure to know how much memory has been allocated to each VM on the host.
Memory in use:	For the control_ domain descriptor, this represents the amount of memory utilized by the control domain. For the xen_server descriptor, this is the host's view of the memory used by the VMs.	MB	A very high value for this measure or a consistent rise in the value of this measure is indicative of abnormal memory usage by the control domain and/or the VMs (depending upon the descriptor). In such circumstances, you may want to compare the value of this measure between the **control_domain** and the **xen_ server** descriptors to isolate the probable cause for the memory drain - is it owing to memory-hungry processes on the control domain? Or is it because of memory-intensive VM processes?
Utilization of memory allocated to VMs:	For the control_ domain descriptor, this represents the percentage of total physical memory of this host (i.e., the value of the Total memory available measure) that is utilized by the control domain. For the xen_server descriptor, this indicates what percentage of the allocated memory (i.e., the value of the Memory allocated to VMs measure) that the VMs on the host have utilized.	Percent	A very high value for this measure or a consistent rise in the value of this measure is indicative of abnormal memory usage by the control domain and/or the VMs (depending upon the descriptor). In such circumstances, you may want to compare the value of this measure between the **control_domain** and the **xen_ server** descriptors to isolate the probable cause for the memory drain - is it owing to memory-hungry processes on the control domain? Or is it because of memory-intensive VM processes? The detailed diagnosis of this measure, which is available only for the **xen_ server** descriptor, reveals how well each VM utilized the allocated memory resources, and thus points you to memory-hungry VMs.

Measurement	Description	Measurement Unit	Interpretation
Memory reclaimed by dynamic memory:	Indicates the amount of memory that is currently reclaimed by the dynamic memory control (DMC) of the XenServer from its VMs.	MB	XenServer DMC (sometimes known as "dynamic memory optimization", "memory over commit" or "memory ballooning") works by automatically adjusting the memory of running VMs, keeping the amount of memory allocated to each VM between specified minimum and maximum memory values or to meet a fixed target, guaranteeing performance and permitting greater density of VMs per server. If DMC is enabled, even when hosts are full, XenServer will attempt to reclaim memory by reducing the memory allocation of running VMs to their defined static values/within specified dynamic ranges. A high value for this measure implies that a large chunk of memory has been reclaimed from the guests. If the value of this measure is abnormally high, it could mean either or all of the following: • The host could be running seriously short of memory, forcing DMC to reclaim the maximum amount of memory from the VMs; • The static memory target/dynamic minimum memory set for the VMs could be very low, allowing a lot of memory to be reclaimed Either way, if too little memory is available to a VM, it could cause that VM to boot slowly. It could also result in poor VM performance and stability issues. To avoid such eventualities, you need to configure DMC in such a way that the VMs always have enough memory to operate smoothly. Using DMC, it is

Measurement	Description	Measurement Unit	Interpretation
			possible to operate a guest virtual machine in one of two modes: • **Target Mode** : The administrator specifies a memory target for the guest.XenServer adjusts the guest's memory allocation to meet the target. Specifying a target is particularly useful in virtual server environments, and in any situation where you know exactly how much memory you want a guest to use. XenServer will adjust the guest's memory allocation to meet the target you specify. • **Dynamic Range Mode** : The administrator specifies a dynamic memory range for the guest; XenServer chooses a target from within the range and adjusts the guest's memory allocation to meet the target. Specifying a dynamic range is particularly useful in virtual desktop environments, and in any situation where you want XenServer to repartition host memory dynamically in response to changing numbers of guests, or changing host memory pressure. XenServer chooses a target from within the range and adjusts the guest's memory allocation to meet the target. Choose a mode wisely and exercise prudence when configuring the corresponding memory values, as they determine how much memory a VM can use for its own operations, and thus governs how well the VM performs.
Maximum memory reclaimed by	Indicates the maximum amount of memory that	MB	A low value is desired for this measure. A high value could indicate either or all of

Measurement	Description	Measurement Unit	Interpretation
dynamic memory:	was reclaimed by DMC from the VMs on the target XenServer.		the following: • The host could be running seriously short of memory, forcing DMC to reclaim the maximum amount of memory from the VMs; • The static memory target/dynamic minimum memory set for one/more VMs could be very low, allowing a lot of memory to be reclaimed Either way, if too little memory is available to a VM, it could cause that VM to boot slowly. It could also result in poor VM performance and stability issues. To avoid such eventualities, you need to configure DMC in such a way that the VMs always have enough memory to operate smoothly.
Utilization of total memory available:	Indicates the percentage of the total memory of the host that is being utilized by the VMs.	Percent	A high value for this measure indicates the existence of one/more resource-hungry VMs on the host. If the condition persists, then the host may not have enough memory to run its own processes. In such situations, you can use the detailed diagnosis of the Utilization of memory allocated to VMs measure to identify the resource-intensive VMs.

The detailed diagnosis of the *Free memory* measure of the **control_domain** descriptor lists the top-10 memory-consuming processes executing on the control domain. Memory-intensive processes can thus be isolated. If required, you can kill one/more of these processes to conserve memory.

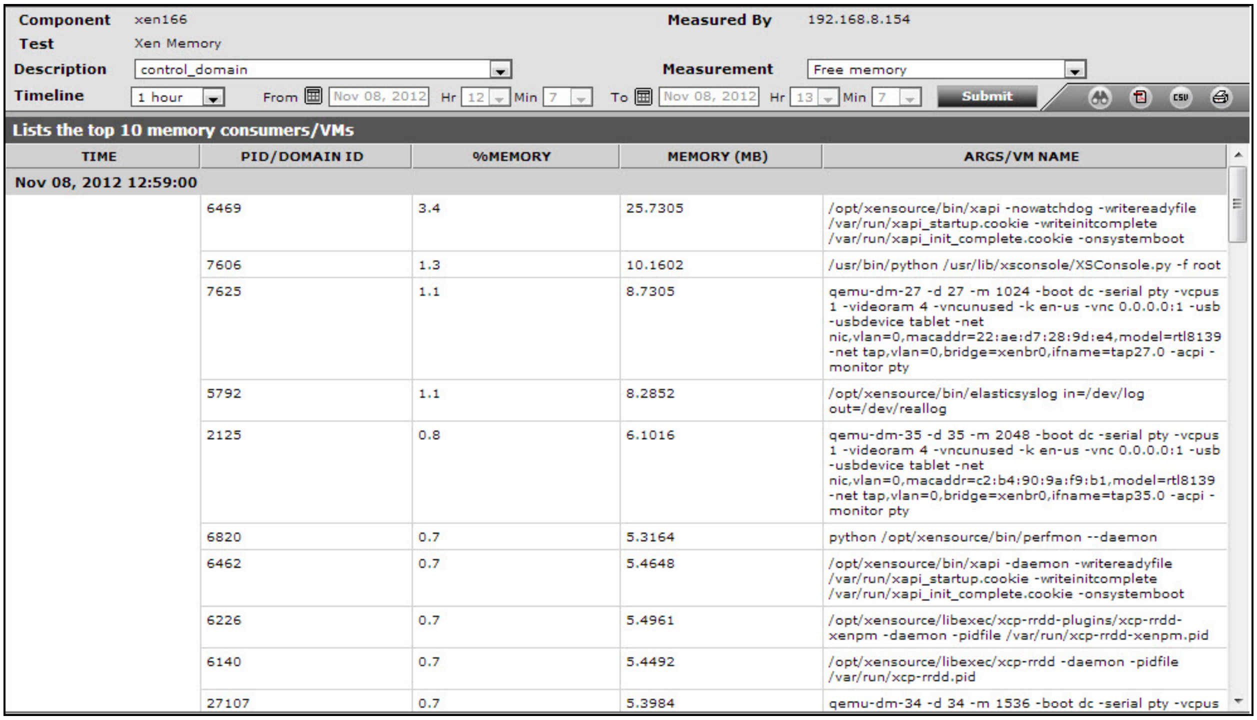

Component xen166 Measured By 192.168.8.154

Test Xen Memory

Description control_domain Measurement Free memory

Timeline 1 hour From Nov 08, 2012 Hr 12 Min 7 To Nov 08, 2012 Hr 13 Min 7 Submit

Lists the top 10 memory consumers/VMs

TIME	PID/DOMAIN ID	%MEMORY	MEMORY (MB)	ARGS/VM NAME
Nov 08, 2012 12:59:00				
	6469	3.4	25.7305	/opt/xensource/bin/xapi -nowatchdog -writereadyfile /var/run/xapi_startup.cookie -writeinitcomplete /var/run/xapi_init_complete.cookie -onsystemboot
	7606	1.3	10.1602	/usr/bin/python /usr/lib/xsconsole/XSConsole.py -f root
	7625	1.1	8.7305	qemu-dm-27 -d 27 -m 1024 -boot dc -serial pty -vcpus 1 -videoram 4 -vncunused -k en-us -vnc 0.0.0.0:1 -usb -usbdevice tablet -net nic,vlan=0,macaddr=22:ae:d7:28:9d:e4,model=rtl8139 -net tap,vlan=0,bridge=xenbr0,ifname=tap27.0 -acpi -monitor pty
	5792	1.1	8.2852	/opt/xensource/bin/elasticsyslog in=/dev/log out=/dev/reallog
	2125	0.8	6.1016	qemu-dm-35 -d 35 -m 2048 -boot dc -serial pty -vcpus 1 -videoram 4 -vncunused -k en-us -vnc 0.0.0.0:1 -usb -usbdevice tablet -net nic,vlan=0,macaddr=c2:b4:90:9a:f9:b1,model=rtl8139 -net tap,vlan=0,bridge=xenbr0,ifname=tap35.0 -acpi -monitor pty
	6820	0.7	5.3164	python /opt/xensource/bin/perfmon --daemon
	6462	0.7	5.4648	/opt/xensource/bin/xapi -daemon -writereadyfile /var/run/xapi_startup.cookie -writeinitcomplete /var/run/xapi_init_complete.cookie -onsystemboot
	6226	0.7	5.4961	/opt/xensource/libexec/xcp-rrdd-plugins/xcp-rrdd-xenpm -daemon -pidfile /var/run/xcp-rrdd-xenpm.pid
	6140	0.7	5.4492	/opt/xensource/libexec/xcp-rrdd -daemon -pidfile /var/run/xcp-rrdd.pid
	27107	0.7	5.3984	qemu-dm-34 -d 34 -m 1536 -boot dc -serial pty -vcpus

Figure 2.8: The detailed diagnosis of the Free memory measure for the control_domain descriptor

The detailed diagnosis of the *Free memory* measure of the **xen_server** descriptor lists the top-10 memory-consuming VMs on the host. Memory-intensive VMs can thus be identified.

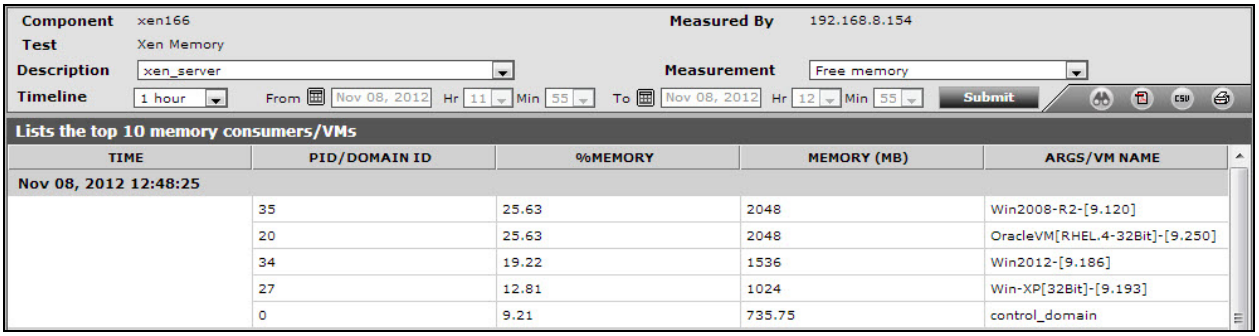

Component xen166 Measured By 192.168.8.154

Test Xen Memory

Description xen_server Measurement Free memory

Timeline 1 hour From Nov 08, 2012 Hr 11 Min 55 To Nov 08, 2012 Hr 12 Min 55 Submit

Lists the top 10 memory consumers/VMs

TIME	PID/DOMAIN ID	%MEMORY	MEMORY (MB)	ARGS/VM NAME
Nov 08, 2012 12:48:25				
	35	25.63	2048	Win2008-R2-[9.120]
	20	25.63	2048	OracleVM[RHEL.4-32Bit]-[9.250]
	34	19.22	1536	Win2012-[9.186]
	27	12.81	1024	Win-XP[32Bit]-[9.193]
	0	9.21	735.75	control_domain

Figure 2.9: The detailed diagnosis of the Free memory measure for the xen_server descriptor

The detailed diagnosis of the *Available memory allocated to VMs* measure indicates how much physical memory has been allocated to each VM on the target host.

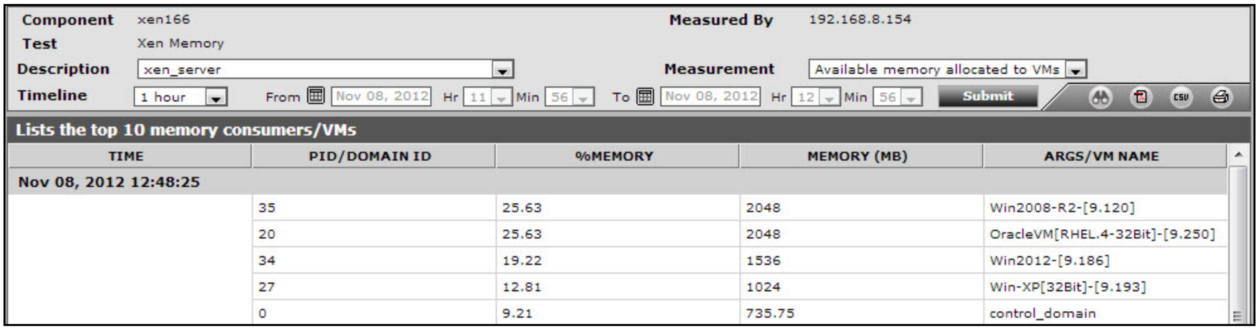

Component xen166 Measured By 192.168.8.154

Test Xen Memory

Description xen_server Measurement Available memory allocated to VMs

Timeline 1 hour From Nov 08, 2012 Hr 11 Min 56 To Nov 08, 2012 Hr 12 Min 56 Submit

Lists the top 10 memory consumers/VMs

TIME	PID/DOMAIN ID	%MEMORY	MEMORY (MB)	ARGS/VM NAME
Nov 08, 2012 12:48:25				
	35	25.63	2048	Win2008-R2-[9.120]
	20	25.63	2048	OracleVM[RHEL.4-32Bit]-[9.250]
	34	19.22	1536	Win2012-[9.186]
	27	12.81	1024	Win-XP[32Bit]-[9.193]
	0	9.21	735.75	control_domain

Figure 2.10: The detailed diagnosis of the Available memory allocated to VMs measure

The detailed diagnosis of the *Memory utilization* measure reveals the percentage of allocated memory that has been utilized by each VM. **Note that the detailed diagnosis capability will be available for the 'xen_ server' descriptor only.** Resource-hungry VMs can thus be identified.

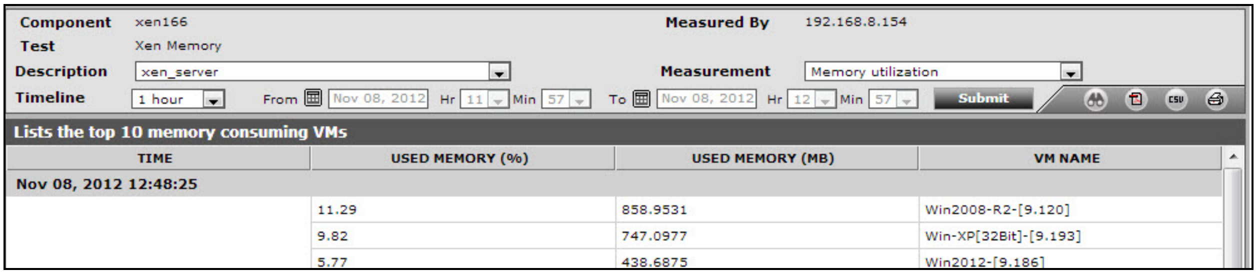

Component xen166 Measured By 192.168.8.154

Test Xen Memory

Description xen_server Measurement Memory utilization

Timeline 1 hour From Nov 08, 2012 Hr 11 Min 57 To Nov 08, 2012 Hr 12 Min 57 Submit

Lists the top 10 memory consuming VMs

TIME	USED MEMORY (%)	USED MEMORY (MB)	VM NAME
Nov 08, 2012 12:48:25			
	11.29	858.9531	Win2008-R2-[9.120]
	9.82	747.0977	Win-XP[32Bit]-[9.193]
	5.77	438.6875	Win2012-[9.186]

Figure 2.11: The detailed diagnosis of the Memory utilization measure

2.1.7 Uptime – Xen Test

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

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

The **Uptime - Xen** test included in the eG agent monitors the uptime of critical XenServers in a virtualized infrastructure.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for the XenServer being monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.

3. **PORT** - The port at which the specified **HOST** listens. By default, this is *NULL*.

4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the **XEN USER** text box. Root user privileges are mandatory when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the root user, then, you have the option of configuring a user with *pool-admin* privileges as the **XEN USER**. If you do not want to expose the credentials of a root/pool-admin user, then you can configure the tests with the credentials of a **XEN USER** with *Read-only* privileges to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the control domain descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a xen user who has *Read-only* privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via **SSH**.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified xen user needs to be mentioned here.

6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.

7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a

non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to 80 or 443 depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

9. **REPORTMANAGERTIME** - By default, this flag is set to **Yes**, indicating that, by default, the detailed diagnosis of this test, if enabled, will report the shutdown and reboot times of the XenServer host in the manager's time zone. If this flag is set to **No**, then the shutdown and reboot times are shown in the time zone of the system where the agent is running (i.e., the system being managed for agent-based monitoring, and the system on which the remote agent is running - for agentless monitoring).

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.

Measurement	Description	Measurement Unit	Interpretation
Has the Xen server been rebooted?	Indicates whether the server has been rebooted during the last measurement period or not.	Boolean	If this measure shows 1, it means that the server 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 server was rebooted. The detailed diagnosis of this measure, if enabled, will provide you with the details of the last reboot of the ESX host. Such details will include the shutdown date/time, reboot date/time, the shutdown duration (in minutes), and whether the host has been configured for maintenance or not.

Measurement	Description	Measurement Unit	Interpretation
Uptime of the Xen server during the last measure period:	Indicates the time period that the system has been up since the last time this test ran.	Secs	If the server 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 server 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 server 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 Xen server:	Indicates the total time that the server has been up since its last reboot.	Mins	Administrators may wish to be alerted if a server 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 **XenUptime** test, then, the **Uptime during the last measure period** measure will report the value 0 until the minute boundary is crossed. For instance, if you configure the XenUptime test host 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 only; 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. This way, every sixth measurement period will report 60 seconds as the uptime of the host. This is because, the XenServer reports uptime only in minutes and not in seconds.

2.1.8 Grid GPUs - Xen Test

GPU-accelerated computing is the use of a graphics processing unit (GPU) together with a CPU to accelerate scientific, analytics, engineering, consumer, and enterprise applications. GPU-accelerated computing enhances application performance by offloading compute-intensive portions of the application to the GPU, while the remainder of the code still runs on the CPU. Architecturally, while a CPU has only few cores and handles few hundred threads at a time, a GPU is composed of hundreds of cores that can handle thousands of threads simultaneously and render a flawless rich graphics experience.

Now, imagine if you could access your GPU-accelerated applications, even those requiring intensive graphics power, anywhere on any device. **NVIDIA GRID** makes this possible. With NVIDIA GRID, a virtualized GPU designed specifically for virtualized server environments, data center managers can bring true PC graphics-rich experiences to users.

The NVIDIA GRID GPUs will be hosted in enterprise data centers and allow users to run virtual desktops or virtual applications on multiple devices connected to the internet and across multiple operating systems, including PCs, notebooks, tablets and even smartphones. Users can utilize their online-connected devices to enjoy the GPU power remotely.

In VDI/virtualized server environments, the NVIDIA GRID delivers GPU resources to virtual desktops/VMs. This way, graphics can be rendered on a virtual machine's (VM's) host server rather than on a physical end-point device. This technology now makes it possible to use virtual desktop technology to support users accessing graphics intensive workloads. There are two modes of making GPU resources available to virtual desktops:

- **Dedicated GPU or GPU Pass-through Technology:** NVIDIA GPU pass-through technology lets you create a virtual workstation that gives users all the benefits of a dedicated graphics processor at their desk. By directly connecting a dedicated GPU to a virtual machine through the hypervisor, you can now allocate the full GPU and graphics memory capability to a single virtual machine without any resource compromise.

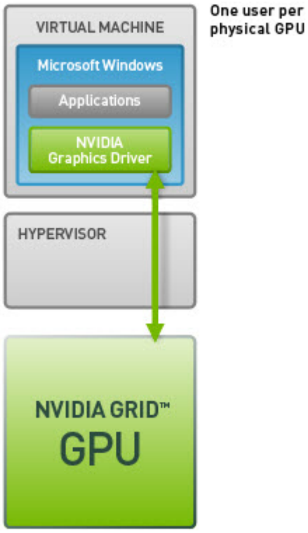

Figure 2.12: Dedicated GPU Technology

- **Shared GPU or Virtual GPU (vGPU) Technology:** GRID vGPU is the industry's most advanced technology for sharing true GPU hardware acceleration between multiple virtual desktops—without compromising the graphics experience. With GRID vGPU technology, the graphics commands of each virtual machine are passed directly to the GPU, without translation by the hypervisor. This allows the GPU hardware to be time-sliced to deliver improved shared virtualized graphics performance. The GRID vGPU manager allows for management of user profiles. IT managers can assign the optimal amount of graphics memory and deliver a customized graphics profile to meet the specific needs of each user. Every virtual desktop has dedicated graphics memory, just like they would at their desk, so they always have the resources they need to launch and run their applications.

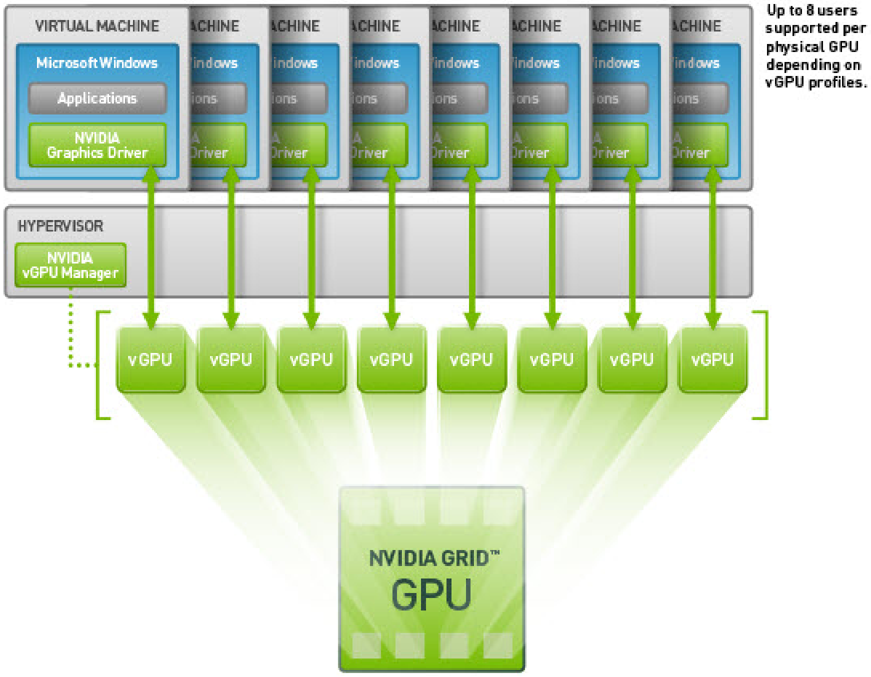

Figure 2.13: Shared vGPU Technology

In GPU-enabled Citrix XenServer environments, if users to VMs/virtual desktops complain of slowness when accessing graphic applications, administrators must be able to instantly figure out what is causing the slowness – is it because adequate GPU resources are not available to the host? Or is it because of excessive utilization of GPU memory and processing resources by a few VMs/virtual desktops on the host? Accurate answers to these questions can help administrators determine whether/not:

- The host is sized with sufficient GPU resources;
- The GPUs are configured with enough graphics and BAR1 memory;
- The GPU technology in use – i.e., the GPU Pass-through technology or the Shared vGPU technology – is ideal for the graphics processing requirements of the environment;

Measures to right-size the host and fine-tune its GPU configuration can be initiated based on the results of this analysis. This is exactly what the **Grid GPUs - Xen** test helps administrators achieve!

This test supports GPU monitoring for NVIDIA K1 and K2 Grids installed on a Citrix XenServer. Using the test, administrators can monitor each physical GPU card installed on the server and can determine how actively memory on that card is utilized, thus revealing the card on which memory is used consistently. In addition, the test also indicates how busy each GPU card is, and in the process pinpoints those GPU cards that are being over-utilized by the VMs/virtual desktops on the host. The adequacy of the physical GPU resources is thus revealed. Moreover, detailed diagnostics provided by the test also lead you to those VMs/virtual desktops that are using each card. In addition, the power consumption and temperature of each

GPU card is also monitored, so that its current temperature and power usage can be ascertained; administrators are thus alerted to abnormal power usage of the GPU and unexpected fluctuations in its temperature. The power limit set and the clock frequencies configured are also revealed, so that administrators can figure out whether the GPU is rightly configured for optimal processing or is any fine-tuning required.

Note:

This test will run and report GPU usage metrics only if the **Shared GPU or Virtual GPU (vGPU) Technology** is used to deliver GPU resources to virtual desktops/VMs. Where the **Dedicated GPU or Pass-through GPU Technology** is used, this test will report only the Virtual machines and the Mode metrics.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each GRID physical GPU assigned to the XenServer host being monitored. Additionally, measures are reported for the summary descriptor too. For the summary descriptor, this test will report the values aggregated across shared GPUs and those that are configured as Unclassified.

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is *NULL*.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the **XEN USER** text box. Root user privileges are mandatory when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the root user, then, you have the option of configuring a user with *pool-admin* privileges as the **XEN USER**. If you do not want to expose the credentials of a root/pool-admin user, then you can configure the tests with the credentials of a **XEN USER** with *Read-only* privileges to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the control domain descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a xen user who has *Read-only* privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via **SSH**.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified xen user needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.
8. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to 80 or 443 depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
GPU memory utilization:	Indicates the proportion of time over the past sample period during which global (device) memory was being read or written on this GPU.	Percent	A value close to 100% is a cause for concern as it indicates that graphics memory on a GPU is almost always in use. In a Shared vGPU environment on the other hand, memory may be consumed all the time, if one/more VMs/virtual desktops on the host utilize the graphics memory excessively and constantly. If you find that only a single VM/virtual desktop has been consistently hogging the graphic memory resources, you may want to switch to the

Measurement	Description	Measurement Unit	Interpretation
			Dedicated GPU mode, so that excessive memory usage by that VM/virtual desktop has no impact on the performance of other VMs/virtual desktops on that host. If the value of this measure is high almost all the time for most of the GPUs, it could mean that the host is not sized with adequate graphics memory.
Used frame buffer memory:	Indicates the amount of frame buffer memory on-board this GPU that is being used by the host.	MB	Frame buffer memory refers to the memory used to hold pixel properties such as color, alpha, depth, stencil, mask, etc. Properties like the screen resolution, color level, and refresh speed of the frame buffer can impact graphics performance. Also, if Error- correcting code (ECC) is enabled on a host, the available frame buffer memory may be decreased by several percent. This is because, ECC uses up memory to detect and correct the most common kinds of internal data corruption. Moreover, the driver may also reserve a small amount of memory for internal use, even without active work on the GPU; this too may impact frame buffer memory. For optimal graphics performance therefore, adequate frame buffer memory should be allocated to the host.
Free frame buffer memory:	Indicates the amount of frame buffer memory on-board this GPU that is yet to be used by the host.	MB	
Frame buffer memory utilization:	Indicates the percentage of frame buffer memory on-board	Percent	A value close to 100% is indicative of excessive frame buffer memory usage. Properties like the screen resolution, color

Measurement	Description	Measurement Unit	Interpretation
	this GPU that is being utilized by the host.		level, and refresh speed of the frame buffer can impact graphics performance. Also, if Error- correcting code (ECC) is enabled on a host, the available frame buffer memory may be decreased by several percent. This is because, ECC uses up memory to detect and correct the most common kinds of internal data corruption. Moreover, the driver may also reserve a small amount of memory for internal use, even without active work on the GPU; this too may impact frame buffer memory. For optimal graphics performance therefore, adequate frame buffer memory should be allocated to the host. Use the detailed diagnosis of this measure to know the frame buffer memory usage of each VM on the monitored server. This will point you to those VMs that are draining the frame buffer memory.
GPU compute utilization:	Indicates the proportion of time over the past sample period during which one or more kernels was executing on this GPU.	Percent	A value close to 100% indicates that the GPU is busy processing graphic requests almost all the time. In a Shared vGPU environment on the other hand, a GPU may be in use almost all the time, if one/more VMs/virtual desktops on the host are running highly graphic-intensive applications. If you find that only a single VM/virtual desktop has been consistently hogging the GPU resources, you may want to switch to the Dedicated GPU mode, so that excessive GPU usage by that VM/virtual desktop has no impact on the performance of other VMs/virtual desktops on that host. If all GPUs are found to be busy most of the time, you may want to consider augmenting

Measurement	Description	Measurement Unit	Interpretation
			the GPU resources of the host. Compare the value of this measure across physical GPUs to know which GPU is being used more than the rest.
Power consumption:	Indicates the current power usage of this GPU.	Watts	A very high value is indicative of excessive power usage by the GPU. In such cases, you may want to enable Power management so that the GPU limits power draw under load to fit within a predefined power envelope by manipulating the current performance state.
Core GPU temperature:	Indicates the current temperature of this GPU.	Celsius	Ideally, the value of this measure should be low. A very high value is indicative of abnormal GPU temperature.
Total framebuffer memory:	Indicates the total size of frame buffer memory of this GPU.	MB	Frame buffer memory refers to the memory used to hold pixel properties such as color, alpha, depth, stencil, mask, etc.
Total BAR1 memory:	Indicates the total size of the BAR1 memory of this GPU.	MB	BAR1 is used to map the frame buffer (device memory) so that it can be directly accessed by the CPU or by 3rd party devices (peer-to-peer on the PCIe bus).
Used BAR1 memory:	Indicates the amount of BAR1 memory on this GPU that is currently being used by the host.	MB	For better user experience with graphic applications, enough BAR1 memory should be available to the host.
Free BAR1 memory:	Indicates the amount of BAR1 memory of this GPU that is still to be used by the host.	MB	
BAR1 memory	Indicates the	Percent	A value close to 100% is indicative of

Measurement	Description	Measurement Unit	Interpretation
utilization:	percentage of the total BAR1 memory on this GPU that is currently being utilized by the host.		excessive BAR1 memory usage by the host. For best graphics performance, sufficient BAR1 memory resources should be available to the host.
Power management:	Indicates whether/not power management is enabled for this GPU.		Many NVIDIA graphics cards support multiple performance levels so that the server can save power when full graphics performance is not required. The default Power Management Mode of the graphics card is Adaptive. In this mode, the graphics card monitors GPU usage and seamlessly switches between modes based on the performance demands of the application. This allows the GPU to always use the minimum amount of power required to run a given application. This mode is recommended by NVIDIA for best overall balance of power and performance. If the power management mode is set to Adaptive, the value of this measure will be Supported. Alternatively, you can set the Power Management Mode to Maximum Performance. This mode allows users to maintain the card at its maximum performance level when 3D applications are running regardless of GPU usage. If the power management mode of a GPU is Maximum Performance, then the value of this measure will be Maximum. The numeric values that correspond to these measure values are discussed in the table below:

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td><table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Supported</td><td>1</td></tr><tr><td>Maximum</td><td>0</td></tr></table>
<b>Note:</b>
By default, this measure will report the <b>Measure Value</b>s listed in the table above to indicate the power management status. In the graph of this measure however, the same is represented using the numeric equivalents only.</td></tr>
<tr><td><b>Power limit:</b></td><td>Indicates the power limit configured for this GPU.</td><td>Watts</td><td><b>This measure will report a value only if the value of the ‘Power management’ measure is ‘Supported’.</b>
The power limit setting controls how much voltage a GPU can use when under load. Its not advisable to set the power limit at its maximum – i.e., the value of this measure should not be the same as the value of the Max power limit measure - as it can cause the GPU to behave strangely under duress.</td></tr>
<tr><td><b>Default power limit:</b></td><td>Indicates the default power management algorithm's power ceiling for this GPU.</td><td>Watts</td><td><b>This measure will report a value only if the value of the ‘Power management’ measure is ‘Supported’</b>.</td></tr>
<tr><td><b>Enforced power limit:</b></td><td>Indicates the power management algorithm's power ceiling for this GPU.</td><td>Watts</td><td><b>This measure will report a value only if the value of the ‘Power management’ measure is ‘Supported’.</b>
The total board power draw is manipulated by the power management algorithm such that it stays under the value reported by this measure.</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
Min power limit:	The minimum value that the power limit of this GPU can be set to.	Watts	**This measure will report a value only if the value of the 'Power management' measure is 'Supported'.**
Max power limit:	The maximum value that the power limit of this GPU can be set to.	Watts	**This measure will report a value only if the value of the 'Power management' measure is 'Supported'.** If the value of this measure is the same as that of the Power limit measure, then the GPU may behave strangely.
Graphics clock:	Indicates the current frequency of the graphics clock of this GPU.	MHz	GPU has many more cores than your average CPU but these cores are much simpler and much smaller so that many more actually fit on a small piece of silicon. These smaller, simpler cores go by different names depending upon the tasks they perform. Stream processors are the cores that perform a single thread at a slow rate. But since GPUs contain numerous stream processors, they make overall computation high. The streaming multiprocessor clock refers to how fast the stream processors run. The Graphics clock is the speed at which the GPU operates. The memory clock is how fast the memory on the card runs. By correlating the frequencies of these clocks (i.e., the value of these measures) with the memory usage, power usage, and overall performance of the GPU, you can figure out if overclocking is required or not. Overclocking is the process of forcing a GPU core/memory to run faster than its manufactured frequency. Overclocking can have both positive and negative effects on

Measurement	Description	Measurement Unit	Interpretation
			GPU performance. For instance, memory overclocking helps on cards with low memory bandwidth, and with games with a lot of post-processing/textures/filters like AA that are VRAM intensive. On the other hand, speeding up the operation frequency of a shader/streaming processor/memory, without properly analyzing its need and its effects, may increase its thermal output in a linear fashion. At the same time, boosting voltages will cause the generated heat to sky rocket. If improperly managed, these increases in temperature can cause permanent physical damage to the core/memory or even "heat death". Putting an adequate cooling system into place, adjusting the power provided to the GPU, monitoring your results with the right tools and doing the necessary research are all critical steps on the path to safe and successful overclocking.
Streaming multiprocessor clock:	Indicates the current frequency of the streaming multiprocessor clock of this GPU.	MHz	
Memory clock:	Indicates the current frequency of the memory clock of this GPU.	MHz	
Fan speed:	Indicates the percent of maximum speed that this GPU's fan is currently intended to run at.	Percent	The value of this measure could range from 0 to 100%. An abnormally high value for this measure could indicate a problem condition – eg., a sudden surge in the temperature of the GPU that could cause the fan to spin faster. Note that the reported speed is only the intended fan speed. If the fan is physically blocked and unable to spin, this output will not match the actual fan speed. Many parts do not report fan speeds because they rely on cooling via fans in the surrounding enclosure. By default the fan speed is increased or decreased automatically in response to changes in temperature.

Measurement	Description	Measurement Unit	Interpretation
Compute processes:	Indicates the number of processes having compute context on this GPU.	Number	Use the detailed diagnosis of this measure to know which processes are currently using the GPU. The process details provided as part of the detailed diagnosis include, the PID of the process, the process name, and the GPU memory used by the process. Note that the GPU memory usage of the processes will not be available in the detailed diagnosis, if the Windows platform on which XenApp operates is running in the WDDM mode. In this mode, the Windows KMD manages all the memory, and not the NVIDIA driver. Therefore, the NVIDIA SMI commands that the test uses to collect metrics will not be able to capture the GPU memory usage of the processes.
Volatile single bit errors:	Indicates the number of volatile single bit errors in this GPU.	Number	Volatile error counters track the number of errors detected since the last driver load. Single bit ECC errors are automatically corrected by the hardware and do not result in data corruption. Ideally, the value of this measure should be 0.
Volatile double bit errors:	Indicates the total number of volatile double bit errors in this GPU.	Number	Volatile error counters track the number of errors detected since the last driver load. Double bit errors are detected but not corrected. Ideally, the value of this measure should be 0.
Aggregate single bit errors:	Indicates the total number of aggregate single bit errors in this GPU.	Number	Aggregate error counts persist indefinitely and thus act as a lifetime counter. Single bit ECC errors are automatically corrected by the hardware and do not result in data corruption.

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>Ideally, the value of this measure should be 0.</td></tr>
<tr><td><b>Aggregate double bit errors:</b></td><td>Indicates the total number of aggregate double bit errors in this GPU.</td><td>Number</td><td>Aggregate error counts persist indefinitely and thus act as a lifetime counter. Double bit errors are detected but not corrected.
Ideally, the value of this measure should be 0.</td></tr>
<tr><td><b>Virtual machines:</b></td><td>Indicates the number of virtual machines allocated to this GPU Grid card.</td><td>Number</td><td>If the Mode measure reports the value Pass-through, then the value of this measure will be 1. Using the detailed diagnosis of this measure, you can identify the VMs that are using the GPU grid card and the process utilization of each VM.</td></tr>
<tr><td><b>Mode:</b></td><td>Indicates the mode using which the GPU resources were delivered to the VMs.</td><td></td><td>The values that this measure can take and their corresponding numeric values are as follows:
<table>
<tr><th>Measure Value</th><th>Numeric Values</th></tr>
<tr><td>Pass through</td><td>0</td></tr>
<tr><td>Shared</td><td>1</td></tr>
<tr><td>Unavailable (GPU card is not allocated to any VM)</td><td>2</td></tr>
</table>
<b>Note:</b>
By default, this test reports the <b>Measure Value</b>s listed in the table above to indicate the mode of GPU delivery. In the graph of this measure however, the same is represented using the numeric equivalents only.</td></tr>
<tr><td><b>Remaining capacity:</b></td><td>Indicates the number of vGPUs in this GPU grid</td><td>Number</td><td>Ideally, the value of this measure should be high.</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
	card that are yet to be allocated to the VMs on this host.		A value close to 0 could indicate that vGPUs have been over-allocated to the VMs. This could result in a serious GPU-contention on the host.

2.2 The Network Layer

The **Network** layer enables administrators to instantaneously detect network connection failures that could render the XenServer host inaccessible to users, and any unexpected / unusual increase in traffic to or from the XenServer host.

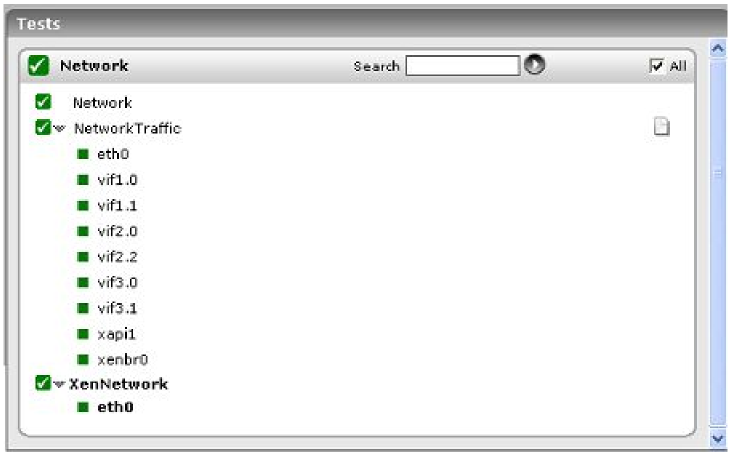

Tests

☑ Network Search ☑ All

☑ Network

☑ NetworkTraffic
- eth0
- vif1.0
- vif1.1
- vif2.0
- vif2.2
- vif3.0
- vif3.1
- xapi1
- xenbr0

☑ XenNetwork
- eth0

Figure 2.14: The tests mapped to the Network layer

The **Network** and **NetworkTraffic** tests have been elaborately discussed in the *Monitoring Unix and Windows Servers* document. The section to come therefore will deal with the XenNetwork test alone.

2.2.1 The Network – Xen Test

This test measures the network I/O activity on the XenServer host.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each network interface supported by a XenServer host

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is *NULL*.

4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the **XEN USER** text box. Root user privileges are mandatory when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the root user, then, you have the option of configuring a user with *pool-admin* privileges as the **XEN USER**. If you do not want to expose the credentials of a root/pool-admin user, then you can configure the tests with the credentials of a **XEN USER** with *Read-only* privileges to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the control domain descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a xen user who has *Read-only* privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via **SSH**.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified xen user needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter

is set to 80 or 443 depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Network data received:	Indicates the rate at which network I/O was read by this network interface.	Mbps	For the **Total** descriptor, this measure will report the total rate at which network I/O was read by the XenServer host across all its network interfaces.
Network data sent:	Indicates the rate at which network I/O was written by this network interface.	Mbps	For the **Total** descriptor, this measure will report the total rate at which network I/O was written by the XenServer host across all its network interfaces.

2.3 The Xen Pool Layer

The test mapped to this layer monitors the resource pool to which the monitored XenServer belongs, and reports whether/not the monitored server is the pool master.

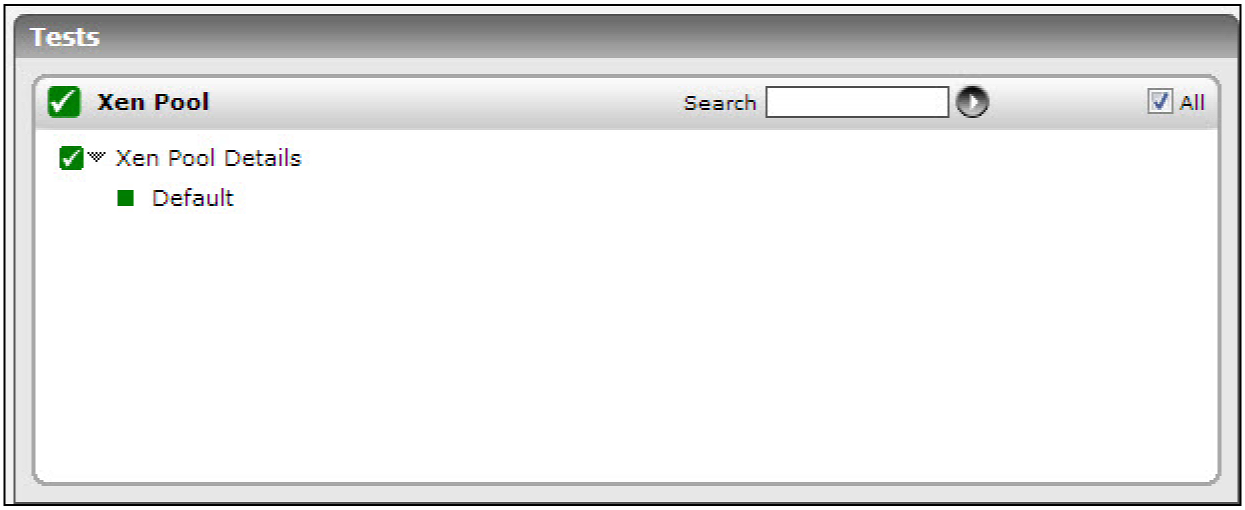

Figure 2.15: The test mapped to the Xen Pool layer

2.3.1 Xen Pool Details Test

A Resource Pool comprises of multiple XenServer Host installations, bound together into a single managed entity which can host Virtual Machines. When combined with shared storage, a Resource Pool enables VMs to be started on any XenServer Host which has sufficient memory and then dynamically moved between XenServer Hosts while running with minimal downtime (XenMotion).

A pool always has at least one physical host, known as the "pool master", that provides a single point of contact for all of the servers in the pool, known as "slaves", managing communication to other members of the pool as necessary. If the pool master is shut down or unavailable, you will not be able to connect to the pool until the master is online again or until you nominate one of the other members as the new pool master for the pool. However, if a pool is High Availability-enabled, then, upon the failure of the master, another host in the pool is automatically selected as the master. VMs in the pool then automatically restart on the new master.

Likewise, you can also enable the Workload balancing component on a pool. Workload Balancing is a XenServer component, packaged as a virtual appliance, that:

a. Creates reports about VM performance in your XenServer environment

b. Evaluates resource utilization and locates virtual machines on the best possible hosts in the pool for their workload's needs

Using this test, you can determine whether/not the XenServer being monitored is the pool master, and if so, understand the composition of the pool and know the status of the hosts in the pool. In addition, for the pool master, this test reports whether/not the HA and Workload balancing features are enabled for the pool.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for the pool to which the monitored XenServer belongs; if the target XenServer is not part of any pool, then this test will report metrics for a **Default** descriptor

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is *NULL*.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the **XEN USER** text box. Root user privileges are mandatory when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the root user, then, you have the option of configuring a user with *pool-admin* privileges as the **XEN USER**. If you do not want to expose the credentials of a root/pool-admin user, then you can configure the tests with the credentials of a **XEN USER** with *Read-only* privileges to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the control domain descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a xen user who has *Read-only* privileges to the

XenServer:

- Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
- In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
- Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via **SSH**.
- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified xen user needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.
8. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to 80 or 443 depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.
9. **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.

Measurements made by the test

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td><b>Is this server pool master?:</b></td><td>Indicates whether/not the monitored XenServer is the master in this pool.</td><td></td><td>If the monitored XenServer is the pool master, then this measure will report the value <i>Yes</i>. If not, then, this measure will report the value <i>No</i>.
The numeric values that correspond to the above-mentioned measure values are as follows:
<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><b>Note:</b>
By default, this test reports the <b>Measures Value</b>s listed in the table above to indicate whether/not a server is the pool master. In the graph of this measure however, the same will be represented using the numeric equivalents.</td></tr>
<tr><td><b>Is this pool high-availability enabled?:</b></td><td>Indicates whether this pool is high-availability (HA) enabled or not?</td><td></td><td><b>This measure is reported only if the XenServer being monitored is the pool master – i.e., only if the ‘Is this server pool master?’ measure reports the value ‘Yes’.</b>
If the pool for which the target XenServer is the master is HA-enabled, then this measure will report the value</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td><p><em>Yes</em>. If not, then, this measure will report the value <em>No</em>.</p><p>The numeric values that correspond to the above-mentioned measure values are as follows:</p><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><p><strong>Note:</strong></p><p>By default, this test reports the <strong>Measures Value</strong>s listed in the table above to indicate whether/not the pool is HA-enabled. In the graph of this measure however, the same will be represented using the numeric equivalents.</p></td></tr>
<tr><td><strong>High availability host failures to be tolerated:</strong></td><td>Indicates the number of failures that this host can tolerate before the pool is declared to be overcommitted.</td><td>Number</td><td><p><strong>This measure is reported only if the XenServer being monitored is the pool master – i.e., only if the ‘Is this server pool master?’ measure reports the value ‘Yes’.</strong></p><p>High Availability works by creating a failure plan (that is, by calculating how many hosts can be restarted based on the priorities you set). The number of hosts that can be restarted is based on the available resources (CPU, memory) in the pool. As you specify the restart priority for VMs, XenServer evaluates the resources required to start each VM. When there are not enough resources to restart all the VMs set to be restarted, the pool reaches its Maximum failure capacity and is considered overcommitted. The pool can also be</p></td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>overcommitted for reasons such as not enough free memory or changes to virtual disks and networks that affect which VMs can be restarted on which servers.
To increase the maximum failure capacity for a pool, you need to do one or more of the following:
• Reduce the number of VMs set to Restart as their restart priority.
• Increase the amount of RAM on your servers or add more servers to the pool to increase its capacity.
• Reduce the amount of memory configured on some VMs.
• Shut down non-essential VMs.</td></tr>
<tr><td><b>Is this pool work load balancing enabled?</b></td><td>Indicates whether/not this pool is enabled for workload balancing.</td><td></td><td><b>This measure is reported only if the XenServer being monitored is the pool master – i.e., only if the ‘Is this server pool master?’ measure reports the value ‘Yes’.</b>
If the pool for which the target XenServer is the master is workload balancing-enabled, then this measure will report the value <i>Yes</i>. If not, then, this measure will report the value <i>No</i>.
The numeric values that correspond to the above-mentioned measure values are as follows:
<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>
<b>Note:</b></td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			By default, this test reports the **Measures Value**s listed in the table above to indicate whether/not the pool is workload balancing- enabled. In the graph of this measure however, the same will be represented using the numeric equivalents.
Total hosts in pool:	Indicates the number of XenServer hosts in this pool.	Number	This measure is reported only if the XenServer being monitored is the pool master – i.e., only if the 'Is this server pool master?' measure reports the value 'Yes'.
Online hosts in pool:	Indicates the number of XenServer hosts in this pool that are currently online.	Number	**This measure is reported only if the XenServer being monitored is the pool master – i.e., only if the 'Is this server pool master?' measure reports the value 'Yes'.**
Offline hosts in pool:	Indicates the number of XenServer hosts in this pool that are currently offline.	Number	**This measure is reported only if the XenServer being monitored is the pool master – i.e., only if the 'Is this server pool master?' measure reports the value 'Yes'.** If the value of this measure is equal to the value of the *Total hosts in pool* measure, it indicates that none of the hosts in the pool are currently available. In this situation, users will neither be able to access the pool or its VMs. If the pool master of a pool that is not HA- enabled goes offline, the slaves realize that communication has been lost and each retry for sixty seconds.

Measurement	Description	Measurement Unit	Interpretation
			Each slave then puts itself into emergency mode, whereby the slave hosts will now only accept the pool emergency commands. If the master comes back up at this point, it will reestablish communication with its slaves, they will leave emergency mode, and operation will return to normal. If the master remains offline, you should choose a slave and promote it to master. Once a slave becomes the master, you need to inform the other slaves who the new master is. Until this process is complete, you will not be able to access the pool. Now, if the slaves in a pool that is not HA-enabled go offline, they will stop sending heartbeat messages to the master. If no heartbeat has been received for 30 seconds then the master assumes the slave is dead. To recover from this problem, you can repair the slave or instruct the master to forget about the slave node. In the case of the latter, all VMs running on the slave will be marked as 'offline' and can be restarted on other hosts. In case of HA-enabled pools, if any host (be it the master or a slave) in the pool goes offline, the HA mechanism automatically moves protected VMs to a healthy host. Additionally, if the host that fails is the master, HA selects another host to take over the master role automatically, meaning that you can continue to manage the XenServer pool.
Disabled hosts in pool:	Indicates the number of XenServer hosts in this pool	Number	**This measure is reported only if the XenServer being monitored is the**

Measurement	Description	Measurement Unit	Interpretation
	that are currently disabled.		**pool master – i.e., only if the 'Is this server pool master?' measure reports the value 'Yes'.** If the server in a resource pool is placed in the Maintenance mode, then all running VMs will be automatically migrated from it to another server in the same pool. If the server is the pool master, a new master will also be selected for the pool. When all running VMs have been successfully migrated off the server, the server's status is changed to and set to Disabled.

2.4 The Outside View of VMs Layer

The **Outside View of VMs** 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 Xen server, so as to minimize the impact it has on the other guests on the current server.
- Determine times when sudden or steady spikes in the physical resource utilization are caused by the guest machines

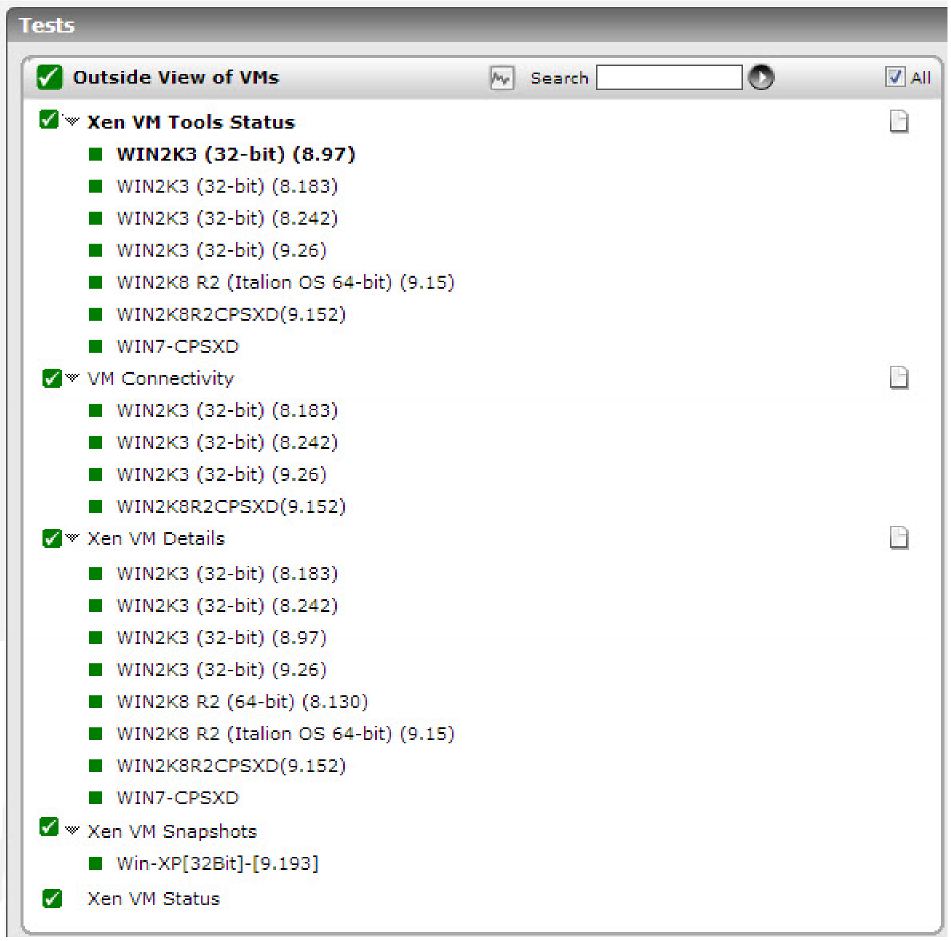

Tests

☑ Outside View of VMs Search ☑ All

☑ Xen VM Tools Status
- 🟩 **WIN2K3 (32-bit) (8.97)**
- 🟩 WIN2K3 (32-bit) (8.183)
- 🟩 WIN2K3 (32-bit) (8.242)
- 🟩 WIN2K3 (32-bit) (9.26)
- 🟩 WIN2K8 R2 (Italion OS 64-bit) (9.15)
- 🟩 WIN2K8R2CPSXD(9.152)
- 🟩 WIN7-CPSXD

☑ VM Connectivity
- 🟩 WIN2K3 (32-bit) (8.183)
- 🟩 WIN2K3 (32-bit) (8.242)
- 🟩 WIN2K3 (32-bit) (9.26)
- 🟩 WIN2K8R2CPSXD(9.152)

☑ Xen VM Details
- 🟩 WIN2K3 (32-bit) (8.183)
- 🟩 WIN2K3 (32-bit) (8.242)
- 🟩 WIN2K3 (32-bit) (8.97)
- 🟩 WIN2K3 (32-bit) (9.26)
- 🟩 WIN2K8 R2 (64-bit) (8.130)
- 🟩 WIN2K8 R2 (Italion OS 64-bit) (9.15)
- 🟩 WIN2K8R2CPSXD(9.152)
- 🟩 WIN7-CPSXD

☑ Xen VM Snapshots
- 🟩 Win-XP[32Bit]-[9.193]

☑ Xen VM Status

Figure 2.16: The tests mapped to the Outside View of VMs layer

2.4.1 VM Connectivity Test

Sometimes, a VM could be in a powered-on state, but the failure of the VM operating system or any fatal error in VM operations could have rendered the VM inaccessible to users. In order to enable administrators to promptly detect such ‘hidden’ anomalies, the eG agent periodically runs a connectivity check on each VM using the VM Connectivity test, and reports whether the VM is accessible over the network or not.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each VM configured on the XenServer host being monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.

4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter

is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

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 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. **PACKETSIZE** - The size of packets used for the test (in bytes)
11. **PACKETCOUNT** - The number of packets to be transmitted during the test
12. **TIMEOUT** - How long after transmission should a packet be deemed lost (in seconds)
13. **PACKETINTERVAL** - Represents the interval (in milliseconds) between successive packet transmissions during the execution of the network test for a specific target.
14. **REPORTUNAVAILABILITY** – By default, this flag is set to **No**. This implies that, by default, the test will not report the unavailability of network connection to any VM. In other words, if the *Network availability of VM* measure of this test registers the value *0*for any VM, then, by default, this test will not report any measure for that VM; under such circumstances, the corresponding VM name will not appear as a descriptor of this test. You can set this flag to **Yes**, if you want the test to report and alert you to the unavailability of the network connection to a VM.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Avg network delay:	Indicates the average delay between transmission of packet to a VM and receipt of the response to the packet at the source.	Secs	An increase in network latency could result from misconfiguration of the router (s) along the path, network congestion, retransmissions at the network, etc.
Min network delay:	The minimum time between	Secs	A significant increase in the minimum

Measurement	Description	Measurement Unit	Interpretation
	transmission of a packet and receipt of the response back.		round-trip time is often a sure sign of network congestion.
Packet loss:	Indicates the percentage of packets lost during transmission from source to target and back.	Percent	Packet loss is often caused by network buffer overflows at a network router or by packet corruptions over the network. The detailed diagnosis for this measure provides a listing of routers that are on the path from the external agent to target server, and the delays on each hop. This information can be used to diagnose the hop(s) that could be causing excessive packet loss/delays.
Network availability of VM:	Indicates whether the network connection is available or not.	Percent	A value of 100 indicates that the VM is connected. The value 0 indicates that the VM is not connected. Typically, the value 100 corresponds to a Packet loss of 0.

2.4.2 VM Status – Xen Test

This test enables administrators to determine how many guests have registered with the XenServer, and how many of these are currently running. In addition, the test also indicates whether any guests have migrated to or from the virtual server. XenMotion is a feature using which virtual machines can be moved from server to server within a Citrix XenServer resource pool without service interruption, making zero-downtime server maintenance possible, and enabling administrators to move resource-hungry running applications to take advantage of available compute power. In short, the Xen VM Status test monitors the effectiveness of XenMotion.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for the XenServer being monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed

2. **HOST** - The host for which the test is to be configured.

3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.

4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.

6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.

7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

   - **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

14. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to *1:1*, it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.

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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Added guests:	Indicates the number of guests that were newly added to the XenServer during this measurement period.	Number	The detailed diagnosis of these measures, if enabled, lists the virtual machines that were migrated by XenMotion to or from (as the case may be) the XenServer.
Removed guests:	Indicates the number of guests that were newly removed from the XenServer during this measurement period.	Number	
Registered guests:	Indicates the total number of virtual machines that have been registered with the XenServer.	Number	
Running guests:	Indicates the total number of virtual machines that are currently running.	Number	
Halted guests:	Indicates the number of	Number	

Measurement	Description	Measurement Unit	Interpretation
	guests that were stopped.		
Suspended guests:	Indicates the total number of virtual machines that have been suspended.	Number	You can suspend a XenVM, saving its state to a file. This allows you to shut down the physical XenServer host and later, after rebooting it, return the XenVM to its original running state.

The detailed diagnosis of the *Added Guests* and *Removed guests* measures, if enabled, lists the virtual machines that were migrated by XenMotion to or from (as the case may be) the XenServer (see Figure 2.17).

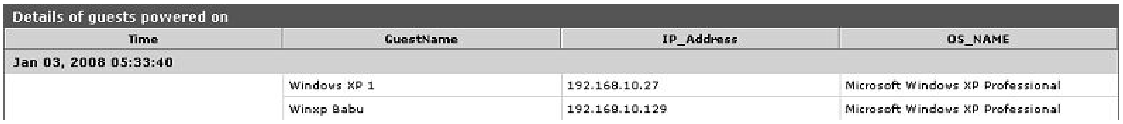

Details of guests powered on

Time	GuestName	IP_Address	OS_NAME
Jan 03, 2008 05:33:40			
	Windows XP 1	192.168.10.27	Microsoft Windows XP Professional
	Winxp Babu	192.168.10.129	Microsoft Windows XP Professional

Figure 2.17: Figure 2.13: The detailed diagnosis of the Added guests measure

The detailed diagnosis of the **Halted guests** measure, if enabled, provides the IP addresses of the guests that stopped running, the guest names, and the operating system on which the guest is executing (see Figure 2.18).

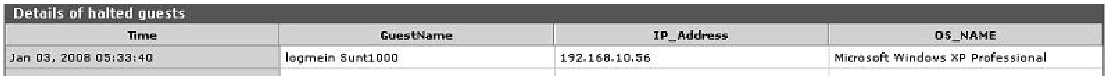

Details of halted guests

Time	GuestName	IP_Address	OS_NAME
Jan 03, 2008 05:33:40	logmein Sunt1000	192.168.10.56	Microsoft Windows XP Professional

Figure 2.18: Figure 2.14: The detailed diagnosis of the Halted guests measure

2.4.3 VM Details – Xen Test

This test monitors the amount of the physical server's resources that each guest on an XenServer 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 taking up the maximum memory utilization, which guest has the maximum disk activity, 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. Also, virtual guests can be organized into resource pools, and allocation of resources can be made at the resource-pool level. In this case, virtual guests allocated to the same resource pool contend for the resources allocated to the resource pool.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each guest configured on the XenServer host

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **REPORT POWERED ON** - You can set the **REPORT POWERED ON** status to **Yes**, so that the test reports an additional measure, Is VM powered on?, revealing whether a guest OS is currently running or not. The default status of this flag is set to Yes for a Citrix XenServer. For a Citrix XenServer – VDI component on the other hand, this flag is set to No by default. This is because, in such environments, the virtual desktops will be in the powered-off state most of the time.

16. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports *80* or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

17. **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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Is VM powered on? :	Whether the virtual machine is currently running on the XenServer host or no.		While the test reports a wide variety of other metrics too for virtual machines

Measurement	Description	Measurement Unit	Interpretation
			that are alive, only the powered on status is indicated for virtual machines that are currently not available. The value *Yes* for this measure indicates that the guest is up and running. The value *No* could indicate that the guest has been powered-off; it could also indicate that XenMotion has moved the guest to a different server. The numeric values that correspond to each of the powered- on states discussed above are listed in the table below: **State** / **Value**: Yes = 1; No = 0 **Note:** By default, this measure reports *Yes* or *No* to indicate VM status. The graph of this measure however, represents the status of a VM using the numeric equivalents - 0 or 1.
Physical CPU usage:	Indicates the percentage of physical CPU used by the guest.	Percent	A high value for this measure indicates a virtual machine that is using a lot of the processor - possibly because one or more processes on this VM are taking a lot of CPU.
Free physical memory:	Indicates the amount of memory available for use with the guest.	MB	Ideally, this value should be high. A low or consistent decrease in this value denotes that the application (s) executing on the guest are consuming memory excessively. You might want to consider increasing the memory allocated to the guest. XenServer

Measurement	Description	Measurement Unit	Interpretation
			Enterprise and XenServer Standard allow that a Linux/Windows VM can use up to 32GB of memory. Moreover, Xen has implemented a balloon driver concept for each domain, enabled independently, that allows the operating system to adjust its current memory allocation up to the maximum limit configured. This allows "unused" allocation to be consumed in other areas, potentially allowing for stable over-commitment of memory resources. Because of this constantly changing memory allocation, memory is allocated and freed dynamically at a granularity of the page-level.
Total physical memory allocated:	Indicates the amount of physical memory currently allocated to the guest.	MB	
Used physical memory:	Indicates the amount of memory used by the guest.	MB	
Usage of allocated memory:	Indicates the percentage of allocated memory that is being used by the guest.	Percent	High memory consumption over long periods can deplete the free memory on the guest, causing prolonged delays in the execution of the application (s) hosted by the guests.
Disk capacity:	Indicates the total allocated disk space of the guest.	MB	
Disk read rate:	Indicates the rate at which the guest read from the disk.	Kbytes/Sec	
Disk write rate:	Indicates the rate at which the guest wrote data to the disk.	Kbytes/Sec	

Measurement	Description	Measurement Unit	Interpretation
Network data received:	Indicates the network I/O reads performed by the guest.	Mbps	
Network data transmitted:	Indicates the network I/O writes performed by the guest.	Mbps	
Virtual CPU utilization:	Indicates the percentage of allocated CPU resources that this VM is currently using.	Percent	Comparing the value of this measure across VMs will enable you to accurately identify the VMs on which CPU- intensive applications are executing.
Virtual CPUs:	Indicates the number of virtual CPU cores allocated to this VM.	Number	
Disk read and write rate:	Indicates the rate at which read- write requests were processed by this VM.	Kbytes/Sec	Compare the value of this measure across VMs to know on which VM I/O activity was abnormally high.
Network data sent and received:	Indicates the rate at which network I/O is processed by this VM.	Mbps	Compare the value of this measure across VMs to know on which VM network I/O activity was abnormally high.
Total IOPS:	Indicates the rate at which I/O operations are performed by this VM.	Requests/Sec	This measure is a good indicator of the I/O processing capacity of the VM. A high value is hence desired for this measure. A consistent drop in this value could indicate a processing bottleneck. In such a situation, you can compare the value of the *Read operations* and *Write operations* measures of the corresponding VM to figure out where the bottleneck lies – in reading data from the VM? or in writing to the VM?

Measurement	Description	Measurement Unit	Interpretation
Read operations:	Indicates the rate at which this VM services read requests.	Requests/Sec	Ideally, the value of this measure should be high. A steady drop in this value indicates a slowdown in processing read requests. Compare the value of this measure across VMs to know which VM is the slowest in responding to read requests.
Write operations:	Indicates the rate at which this VM services write requests.	Requests/Sec	Ideally, the value of this measure should be high. A steady drop in this value indicates a slowdown in processing write requests. Compare the value of this measure across VMs to know which VM is the slowest in responding to write requests.
Time spent waiting for I/O:	Indicates the percentage of time the host's CPU was waiting for this VM to complete I/O processing.	Percent	A high value for this measure indicates that the VM is taking too long to complete I/O processing. This hints at a probable processing bottleneck with the VM.
Average queue size:	Indicates the average number of I/O requests to this VM that are in queue for processing.	Number	If the value of this measure grows consistently, it indicates that the VM is unable to process requests quickly enough to clear the queue. The VM with the maximum number of queued requests could be experiencing a serious I/O processing bottleneck. To identify this VM, compare the value of this measure across VMs.
Current requests in flight:	Indicates the number of I/O requests to this VM that are currently being processed.	Number	

2.4.4 Xen VM Snapshots Test

XenServer provides a convenient snapshotting mechanism that can take a snapshot of a VM storage and metadata at a given time. Where necessary, IO is temporarily halted while the snapshot is being taken to ensure that a selfconsistent disk image can be captured. Snapshot operations result in a snapshot VM that is similar to a template. The VM snapshot contains all the storage information and VM configuration, including attached VIFs, allowing them to be exported and restored for backup purposes.

Though snapshot files are small in size initially, they will grow as writes are made to the VM's disk files. If the number and size of the snapshot files grow significantly over time, they might end up eroding considerable storage space, thereby choking VM operations. To conserve space, administrators need to continuously track snapshot growth per VM, identify 'heavy-weight' snapshots that may not be of use any longer, and purge them. The **Xen VM Snapshots** test helps administrators achieve the same. While the measures reported by the test capture the snapshot file count per VM and the total size of the snapshot files of a VM, the detailed diagnosis reveals the size of each snapshot, thus enabling administrators to quickly spot those snapshot files that are too large in size.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each VM on the target XenServer host

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.

6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.

7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

9. **AGELIMIT** - By default, 15 days is set as **AGELIMIT**. This implies that the test will report all those snapshots that are more than 15 days old as **Old snapshots**. If required, you can change the **AGELIMIT**.

10. **SIZELIMIT**- By default, 10000 KB is set as the **SIZELIMIT**. This implies that the test will report all those snapshots that have a size more than 10000 KB as **Large snapshots**. If required, you can change the **SIZELIMIT**.

11. **SHOW TOP** – By default, the detailed diagnosis of the *Number of snapshots* measure provides details of all the snapshots on a VM. Likewise, the detailed diagnosis of the *Large snapshot count* and *Aged snapshot count* measures provides the details of all large-sized snapshots and old snapshots, respectively. This is why, the **SHOW TOP** parameter is set to *All* by default. To make sure that the detailed diagnosis of the each of these measures lists, for instance, only the top-5 snapshots in terms of size/age (as the case may be), then, specify **5** against **SHOW TOP**. If this is done, then:

    a. The detailed diagnosis of the *Number of snapshots* and *Aged snapshot count* measures will provide the details of only the top-5 snapshots in terms of age.

    b. The detailed diagnosis of the *Large snapshot count* measure will provide the details of only the top-5 snapshots in terms of size.

    In the same way, you can specify any non-zero value against show top to view only that many top

12. **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.

13. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to *1:1*, it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.

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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Number of snapshots:	Indicates the number of snapshot files of this VM that are currently available.	Number	A number of snapshots of a VM provides administrators with multiple restore points. On the flipside though, a high number of snapshots can also be considered a waste of valuable disk space, especially if many of the snapshots hold less critical, but heavy-weight changes/writes to the disk. To accurately identify those snapshots that are consuming disk space excessively, and to learn when they

Measurement	Description	Measurement Unit	Interpretation
			were created, who their parents are, and their current consistent file system state, use the detailed diagnosis of this measure.
Aged snapshots count:	Indicates the number of snapshots that are of an age over the configured **AGELIMIT**.	Number	Use the detailed diagnosis of this measure to identify the old snapshots, so that you can figure out whether they deserve to be retained or not. While many snapshots provide essential restore points for VMs, many others hold less critical information. The 'less useful' snapshots can be eliminated to save disk space.
Large snapshots count:	Indicates the number of snapshots that are of a size more than the configured **SIZELIMIT**.	Number	Though small in size initially, snapshots can grow with time, but can never grow beyond the original disk file size. If a marked increase is noticed in the value of this measure over time, it could indicate that a number of snapshots are rapidly growing in size. To know which snapshots are growing beyond the size limit set, use the detailed diagnosis of this measure.

To accurately identify those snapshots that are consuming disk space excessively, and to learn when they were created, who their parents are, and their current consistent file system state, use the detailed diagnosis of the *Number of snapshots* measure.

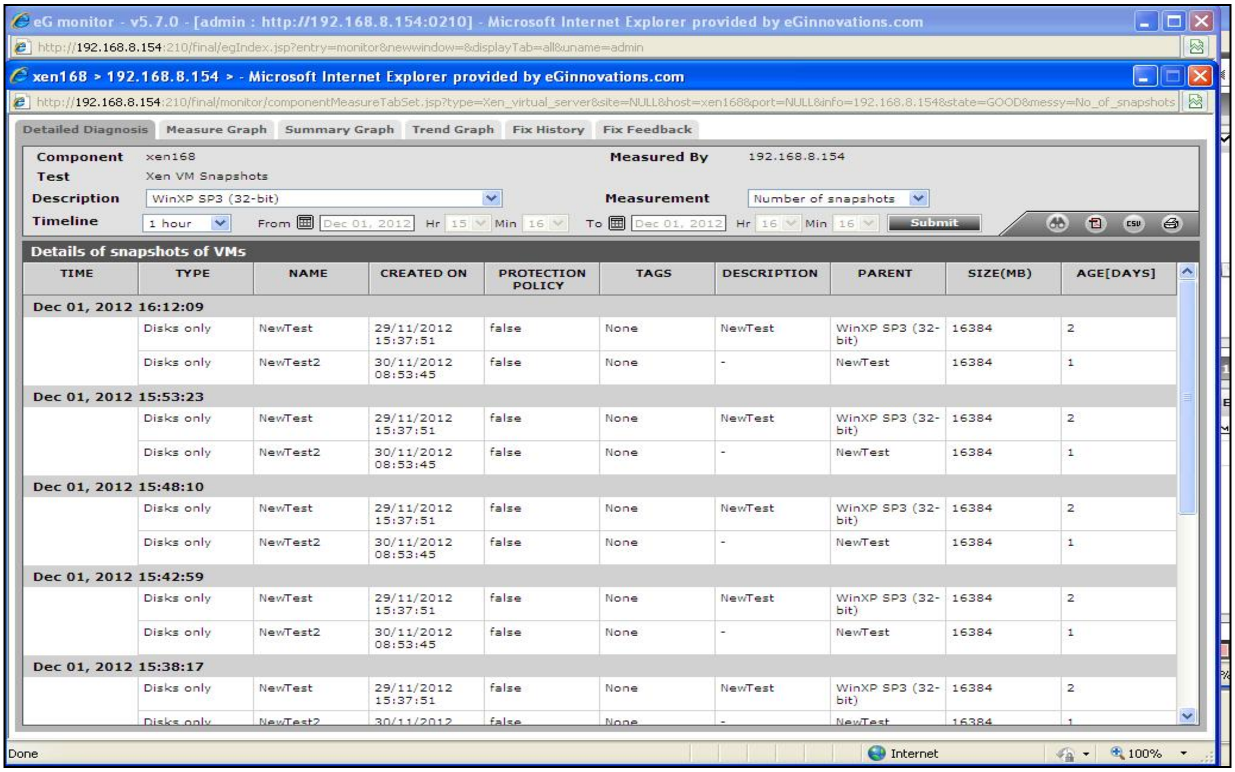

eG monitor - v5.7.0 - [admin : http://192.168.8.154:0210] - Microsoft Internet Explorer provided by eGinnovations.com
http://192.168.8.154:210/final/egIndex.jsp?entry=monitor&newwindow=&displayTab=all&uname=admin
xen168 > 192.168.8.154 > - Microsoft Internet Explorer provided by eGinnovations.com
http://192.168.8.154:210/final/monitor/componentMeasureTabSet.jsp?type=Xen_virtual_server&site=NULL&host=xen168&port=NULL&info=192.168.8.154&state=GOOD&messy=No_of_snapshots
Detailed Diagnosis
Measure Graph
Summary Graph
Trend Graph
Fix History
Fix Feedback
Submit
Done
Internet
100%

Component xen168

Test Xen VM Snapshots

Description WinXP SP3 (32-bit)

Timeline 1 hour From Dec 01, 2012 Hr 15 Min 16 To Dec 01, 2012 Hr 16 Min 16

Measured By 192.168.8.154

Measurement Number of snapshots

Details of snapshots of VMs

TIME	TYPE	NAME	CREATED ON	PROTECTION POLICY	TAGS	DESCRIPTION	PARENT	SIZE(MB)	AGE[DAYS]
Dec 01, 2012 16:12:09									
	Disks only	NewTest	29/11/2012 15:37:51	false	None	NewTest	WinXP SP3 (32-bit)	16384	2
	Disks only	NewTest2	30/11/2012 08:53:45	false	None	-	NewTest	16384	1
Dec 01, 2012 15:53:23									
	Disks only	NewTest	29/11/2012 15:37:51	false	None	NewTest	WinXP SP3 (32-bit)	16384	2
	Disks only	NewTest2	30/11/2012 08:53:45	false	None	-	NewTest	16384	1
Dec 01, 2012 15:48:10									
	Disks only	NewTest	29/11/2012 15:37:51	false	None	NewTest	WinXP SP3 (32-bit)	16384	2
	Disks only	NewTest2	30/11/2012 08:53:45	false	None	-	NewTest	16384	1
Dec 01, 2012 15:42:59									
	Disks only	NewTest	29/11/2012 15:37:51	false	None	NewTest	WinXP SP3 (32-bit)	16384	2
	Disks only	NewTest2	30/11/2012 08:53:45	false	None	-	NewTest	16384	1
Dec 01, 2012 15:38:17									
	Disks only	NewTest	29/11/2012 15:37:51	false	None	NewTest	WinXP SP3 (32-bit)	16384	2
	Disks only	NewTest2	30/11/2012	false	None	-	NewTest	16384	1

Figure 2.19: The detailed diagnosis of the Number of snapshots measure

2.4.5 Xen VM Tools Status Test

XenServer Tools, the XenServer paravirtualization tools, provide high speed I/O for enhanced disk and network performance. XenServer Tools must be installed on each virtual machine (Windows and Linux) in order to use the xe CLI or XenCenter; moreover, VM performance will significantly deteriorate unless the tools are installed.

Without the tools installed, you will not be able to do any the following:

a. Cleanly shut down, reboot or suspend a VM.

b. View VM performance data in XenCenter.

c. Migrate a running VM (via XenMotion).

d. Create quiesced snapshots or snapshots with memory (checkpoints), or revert to snapshots.

e. Adjust the number of vCPUs on a running Linux VM (Windows VMs require a reboot for this to take effect).

Using the **Xen VM Tools Status** test, you can ascertain whether the XenServer tools are installed or not on a VM.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each VM on the target XenServer host

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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.
9. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.
10. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to *1:1*, it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Is Xen tools installed?:	Indicates whether/not XenServers Tools has been		If XenServer Tools is installed on a VM, this measure will report the value *Yes*.

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td>installed on this VM.</td><td></td><td>On the other hand, if XenServer Tools is not installed on a VM, this measure will report the value <i>No</i>.

The numeric values that correspond to each of the measure values discussed above have been detailed 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><b>Note:</b>

By default, this measure reports <i>Yes</i> or <i>No</i> to indicate whether XenServer Tools is installed on a VM or not. The graph of this measure however, represents the same using the numeric equivalents only – i.e., 0 or 1.</td></tr>
<tr><td><b>Is Xen tools upto date?:</b></td><td>Indicates whether the latest version of XenServer Tools has been installed on this VM.</td><td></td><td>If the latest version of XenServer Tools is installed on a VM, this measure will report the value <i>Yes</i>. On the other hand, if an older/obsolete version of XenServer Tools is installed on a VM, this measure will report the value <i>No</i>.

The numeric values that correspond to each of the measure values discussed above have been detailed 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 Yes or</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			No to indicate the version of XenServer Tools installed on a VM. The graph of this measure however, represents the same using the numeric equivalents only – i.e., 0 or 1.

2.5 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 XenServer host is 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 XenServer server host, 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 2.20 appears, which provides you with an overview of individual guest performance.

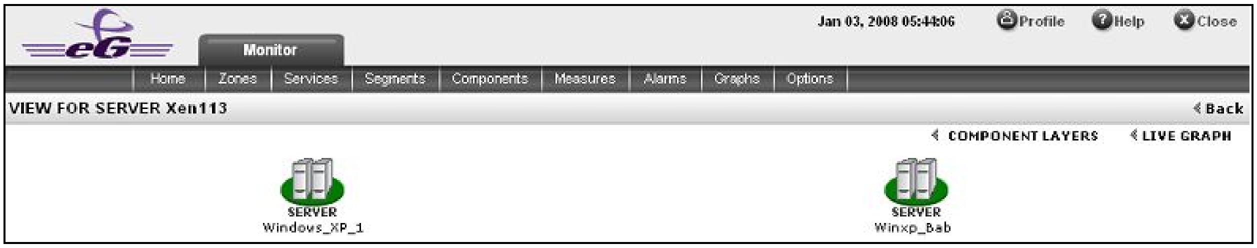

Figure 2.20: A list of guest operating systems on an XenServer host and their current state

To return to the layer model of the *Citrix XenServer* and view the tests associated with the **Inside View of VMs** layer, click on the **COMPONENT LAYERS** link in Figure 2.20. You can now view the list of tests mapped to the **Inside View of VMs** layer, as depicted by Figure 2.21 below.

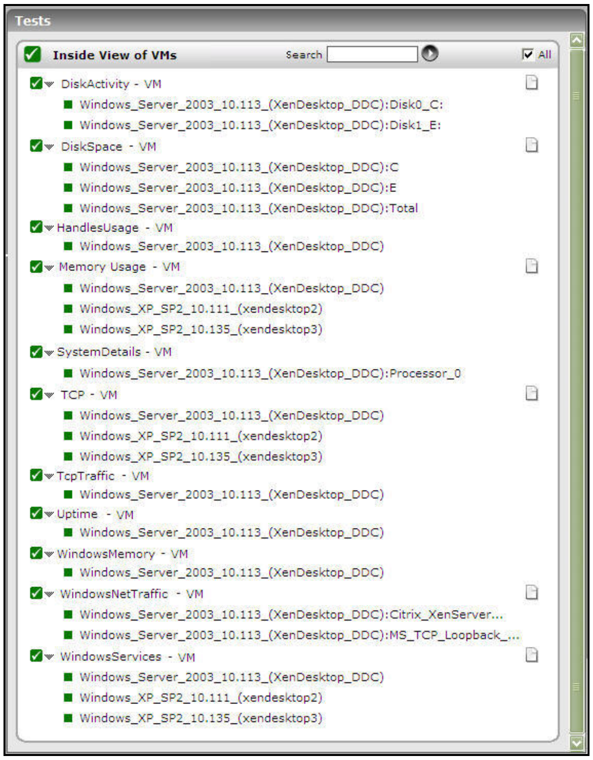

Tests

☑ Inside View of VMs   Search   ☑ All

- ☑ DiskActivity - VM
  - Windows_Server_2003_10.113_(XenDesktop_DDC):Disk0_C:
  - Windows_Server_2003_10.113_(XenDesktop_DDC):Disk1_E:
- ☑ DiskSpace - VM
  - Windows_Server_2003_10.113_(XenDesktop_DDC):C
  - Windows_Server_2003_10.113_(XenDesktop_DDC):E
  - Windows_Server_2003_10.113_(XenDesktop_DDC):Total
- ☑ HandlesUsage - VM
  - Windows_Server_2003_10.113_(XenDesktop_DDC)
- ☑ Memory Usage - VM
  - Windows_Server_2003_10.113_(XenDesktop_DDC)
  - Windows_XP_SP2_10.111_(xendesktop2)
  - Windows_XP_SP2_10.135_(xendesktop3)
- ☑ SystemDetails - VM
  - Windows_Server_2003_10.113_(XenDesktop_DDC):Processor_0
- ☑ TCP - VM
  - Windows_Server_2003_10.113_(XenDesktop_DDC)
  - Windows_XP_SP2_10.111_(xendesktop2)
  - Windows_XP_SP2_10.135_(xendesktop3)
- ☑ TcpTraffic - VM
  - Windows_Server_2003_10.113_(XenDesktop_DDC)
- ☑ Uptime - VM
  - Windows_Server_2003_10.113_(XenDesktop_DDC)
- ☑ WindowsMemory - VM
  - Windows_Server_2003_10.113_(XenDesktop_DDC)
- ☑ WindowsNetTraffic - VM
  - Windows_Server_2003_10.113_(XenDesktop_DDC):Citrix_XenServer...
  - Windows_Server_2003_10.113_(XenDesktop_DDC):MS_TCP_Loopback_...
- ☑ WindowsServices - VM
  - Windows_Server_2003_10.113_(XenDesktop_DDC)
  - Windows_XP_SP2_10.111_(xendesktop2)
  - Windows_XP_SP2_10.135_(xendesktop3)

Figure 2.21: The tests associated with the Virtual Servers layer

As indicated in Figure 2.21, 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 XenServer. 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.

2.5.1 Disk Activity - VM Test

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

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every combination of *virtual_guest:disk_partition*

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the

XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

- Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
- In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
- Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.
8. **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*.

9. **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 **2.7.2**.

   - **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

16. **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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
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	

Measurement	Description	Measurement Unit	Interpretation
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.

Measurement	Description	Measurement Unit	Interpretation
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**.

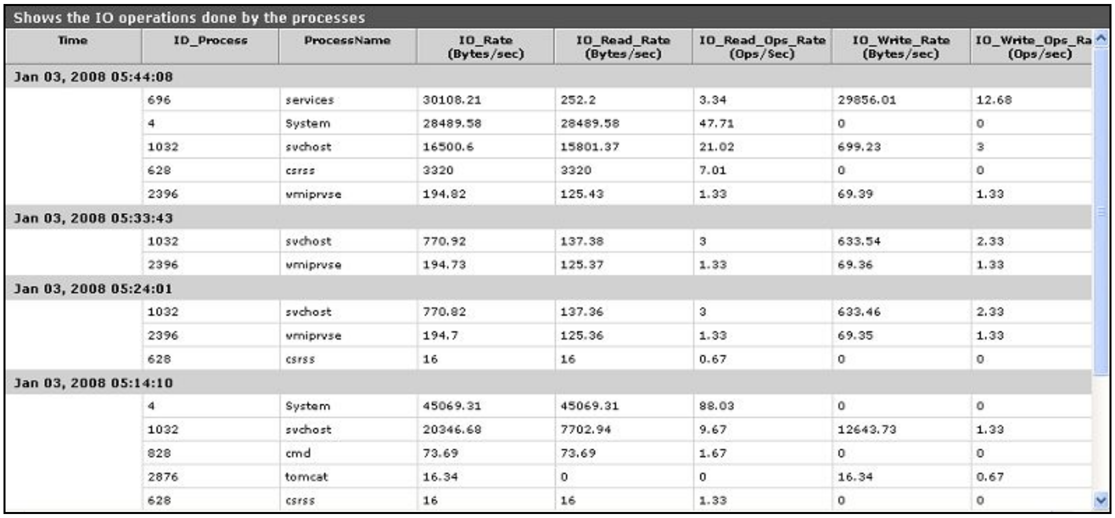

Shows the IO operations done by the processes

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 2.22: The detailed diagnosis of the Percent virtual busy measure

2.5.1.1 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.

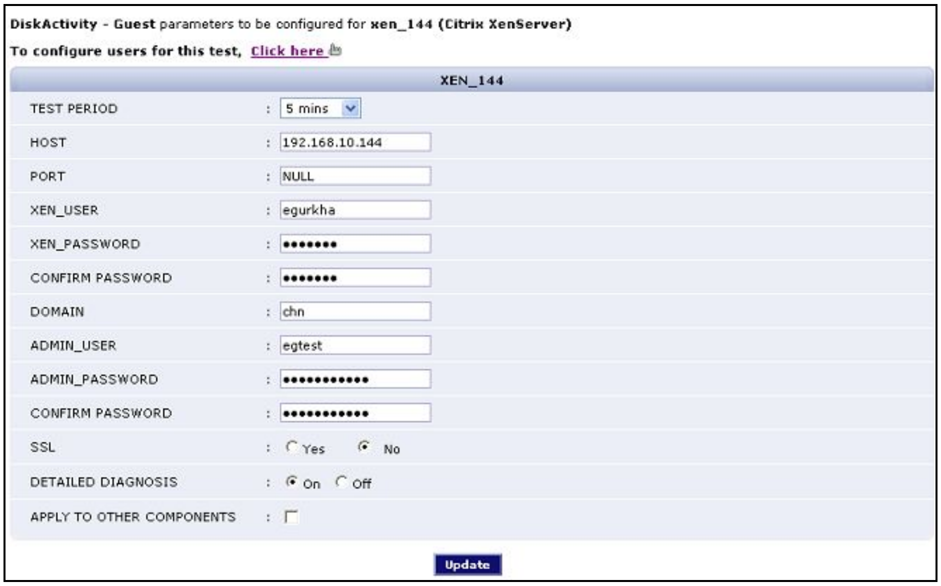

DiskActivity - Guest parameters to be configured for **xen_144 (Citrix XenServer)**

To configure users for this test, Click here

XEN_144

TEST PERIOD	:	5 mins
HOST	:	192.168.10.144
PORT	:	NULL
XEN_USER	:	egurkha
XEN_PASSWORD	:	•••••••
CONFIRM PASSWORD	:	•••••••
DOMAIN	:	chn
ADMIN_USER	:	egtest
ADMIN_PASSWORD	:	•••••••••••
CONFIRM PASSWORD	:	•••••••••••
SSL	:	○ Yes ● No
DETAILED DIAGNOSIS	:	● On ○ Off
APPLY TO OTHER COMPONENTS	:	□

Figure 2.23: Configuring a VM test

Upon clicking, Figure 2.24 will appear, using which the VM user details can be configured.

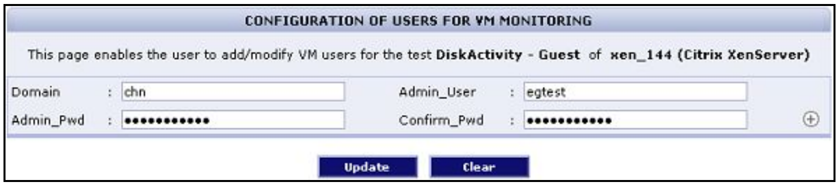

CONFIGURATION OF USERS FOR VM MONITORING

This page enables the user to add/modify VM users for the test DiskActivity - Guest of xen_144 (Citrix XenServer)

Domain	:	chn	Admin_User	:	egtest
Admin_Pwd	:	●●●●●●●●●●●	Confirm_Pwd	:	●●●●●●●●●●●

Figure 2.24: 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 2.24). 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 2.7.2 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 2.24. This will allow you to add one more user specification as depicted by Figure 2.25.

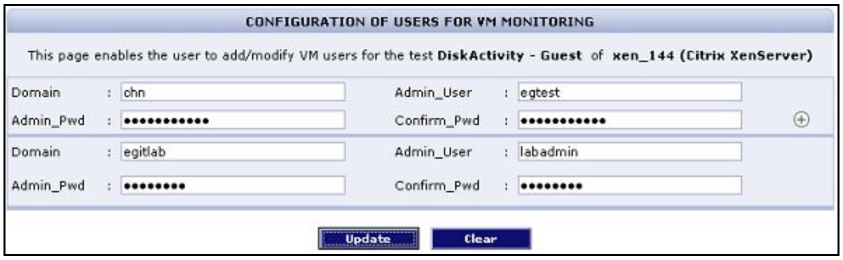

CONFIGURATION OF USERS FOR VM MONITORING

This page enables the user to add/modify VM users for the test DiskActivity - Guest of xen_144 (Citrix XenServer)

Domain	:	chn	Admin_User	:	egtest
Admin_Pwd	:	●●●●●●●●●●●	Confirm_Pwd	:	●●●●●●●●●●●
Domain	:	egitlab	Admin_User	:	labadmin
Admin_Pwd	:	●●●●●●●●	Confirm_Pwd	:	●●●●●●●●

Figure 2.25: 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 2.26.

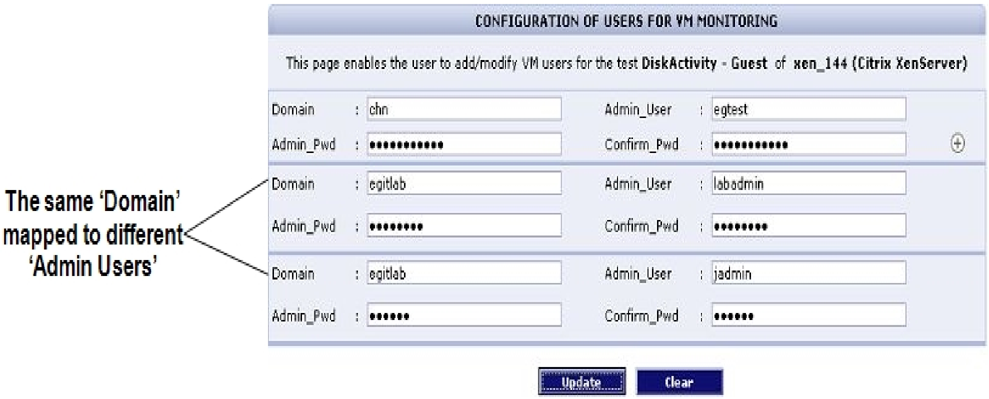

The same 'Domain' mapped to different 'Admin Users'

CONFIGURATION OF USERS FOR VM MONITORING

This page enables the user to add/modify VM users for the test **DiskActivity - Guest** of **xen_144 (Citrix XenServer)**

Domain	:	chn	Admin_User	:	egtest
Admin_Pwd	:	●●●●●●●●●●●	Confirm_Pwd	:	●●●●●●●●●●●
Domain	:	egitlab	Admin_User	:	labadmin
Admin_Pwd	:	●●●●●●●●	Confirm_Pwd	:	●●●●●●●●
Domain	:	egitlab	Admin_User	:	jadmin
Admin_Pwd	:	●●●●●●	Confirm_Pwd	:	●●●●●●

Figure 2.26: 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 2.26, 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 2.26). 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 2.26.

8. To remove the details of a particular user alone, just click the ⊖ button in Figure 2.26.

9. To save the specification, just click on the **Update** button in Figure 2.26. 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 2.26).

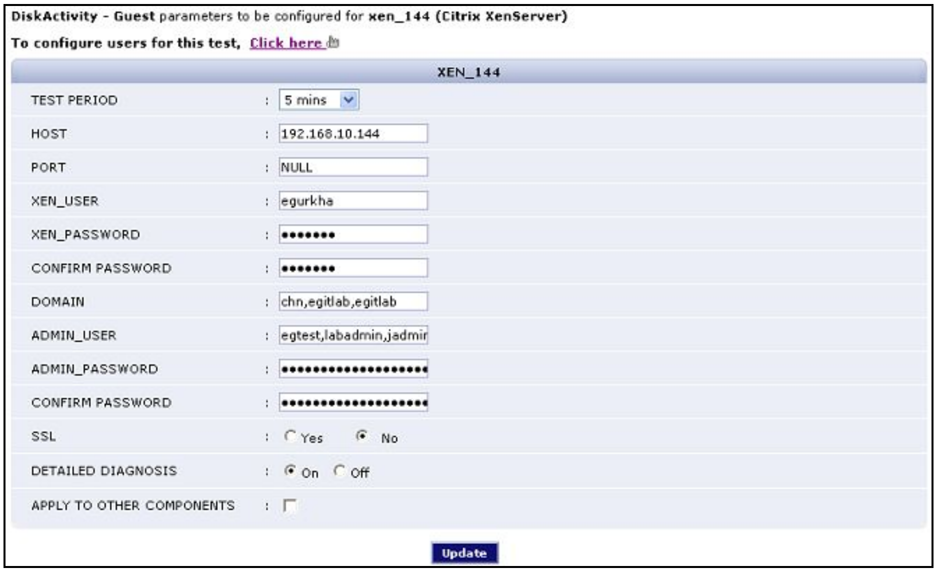

DiskActivity - Guest parameters to be configured for **xen_144 (Citrix XenServer)**

To configure users for this test, Click here

XEN_144

TEST PERIOD	:	5 mins
HOST	:	192.168.10.144
PORT	:	NULL
XEN_USER	:	egurkha
XEN_PASSWORD	:	●●●●●●●
CONFIRM PASSWORD	:	●●●●●●●
DOMAIN	:	chn,egitlab,egitlab
ADMIN_USER	:	egtest,labadmin,jadmir
ADMIN_PASSWORD	:	●●●●●●●●●●●●●●●●●●●
CONFIRM PASSWORD	:	●●●●●●●●●●●●●●●●●●●
SSL	:	○ Yes ● No
DETAILED DIAGNOSIS	:	● On ○ Off
APPLY TO OTHER COMPONENTS	:	□

Figure 2.27: The test configuration page displaying multiple domain names, user names, and passwords

2.5.2 Disk Space - VM Test

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

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every combination of *virtual_guest:disk_partition*

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes

bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

- Obtain the server-certificate for the XenServer
- Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total capacity:	Indicates the total capacity of a disk partition.	MB	
Used space:	Indicates the amount of space used in a disk partition.	MB	
Free space:	Indicates the current free space available for each disk partition of a system.	MB	

Measurement	Description	Measurement Unit	Interpretation
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.

2.5.3 Memory Usage - VM Test

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

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every VM on a Citrix XenServer

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

    - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
    - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
    - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.

- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.

6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.

7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

16. **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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
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

Measurement	Description	Measurement Unit	Interpretation
			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 memoryis Free physical memory. Typically, Available physical memoryis 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

Measurement	Description	Measurement Unit	Interpretation
			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.

Measurement	Description	Measurement Unit	Interpretation
			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.

2.5.4 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:

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every combination of *virtual_guest:disk_partition*

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed

2. **HOST** - The host for which the test is to be configured.

3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.

4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.

   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.

   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.

   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.

6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.

7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer

   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

16. **ENABLE MEMORY DIAGNOSIS** - By default, the **ENABLE MEMORY DIAGNOSIS** flag is set to **NO**, indicating that detailed diagnosis will not be available for the *Free memory in VM* measure reported by this test by default. If you want to view the detailed diagnosis of the *Free memory in VM* measure - i.e., to view the top 10 processes on the target VM that are utilizing memory excessively - you can change this flag to **YES**.

17. **USE TOP FOR DD** - **This parameter is applicable only 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.

18. **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.

Measurements made by the test

Measurement	Description	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.

Measurement	Description	Measurement Unit	Interpretation
			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 2.28). 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.

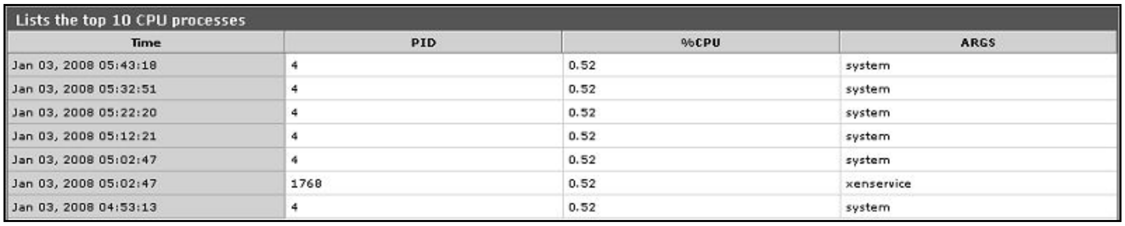

Lists the top 10 CPU processes

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 2.28: 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 2.29). 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.

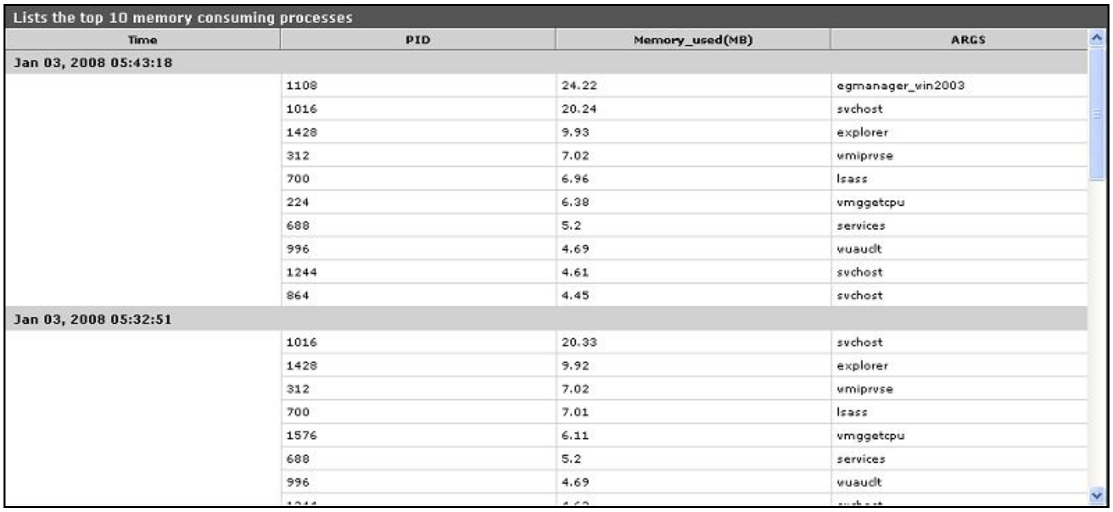

Lists the top 10 memory consuming processes

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	wmiprvse
	700	6.96	lsass
	224	6.38	vmggetcpu
	688	5.2	services
	996	4.69	wuaudt
	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	wmiprvse
	700	7.01	lsass
	1576	6.11	vmggetcpu
	688	5.2	services
	996	4.69	wuaudt

Figure 2.29: The detailed diagnosis of the Free memory measure listing the top 10 memory consuming processes

2.5.5 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.

This test included in the eG agent monitors the uptime of each VM on a XenServer host.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every guest on a Citrix XenServer

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.

3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.

4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.

   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.

   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.

   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.

6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.

7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer

   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

   - **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **REPORTMANAGERTIME** - By default, this flag is set to **Yes**, indicating that, by default, the detailed diagnosis of this test, if enabled, will report the shutdown and reboot times of the VMs in the manager's time zone. If this flag is set to **No**, then the shutdown and reboot times are shown in the time zone of the system where the agent is running (i.e., the system being managed for agent-based monitoring, and the system on which the remote agent is running - for agentless monitoring).

16. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

17. **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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Has the 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.

Measurement	Description	Measurement Unit	Interpretation
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 - VM test, then, the *Uptime during the last measure period* measure will report the value 0 for Unix VMs (only) until the minute boundary is crossed. For instance, if you configure the Uptime - VM 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 Unix 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 Unix VMs. This is because, Unix-based operating systems report uptime only in minutes and not in seconds.

2.5.6 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.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every Windows VM guest/user on the monitored XenServer

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

    - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
    - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
    - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
    - Then, click the **Update** button to save the changes.

    Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent

communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

- Obtain the server-certificate for the XenServer
- Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
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

Measurement	Description	Measurement Unit	Interpretation
			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

Measurement	Description	Measurement Unit	Interpretation
			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.

2.5.7 Windows Network Traffic - VM Test

This is an internal test that monitors the incoming and outgoing traffic through each Windows guest of a XenServer host.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every *Windows_virtual_guest:network_interface* combination

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed

2. **HOST** - The host for which the test is to be configured.

3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.

4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.

   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.

   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.

   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.

6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.

7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer

   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

16. **IS FULL DUPLEX** - By default, this flag is set to **Yes**, indicating that the incoming and outgoing data traffic is handled in full duplex mode. This means that the network interfaces are sending and receiving data at the same time. If the data traffic in your environment is handled in half-duplex mode, set this flag to **No**. This means that the network interfaces are not sending and receiving data at the same time; in essence, it is a one-way conversation. In this case, the test halves the value of the **Incoming traffic** and **Outgoing traffic** measures.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Incoming traffic:	Indicates the rate at which data (including framing characters) is received on a network interface.	Mbps	An abnormally high rate of incoming traffic may require additional analysis.
Outgoing traffic:	Represents the rate at which data (including framing characters) is sent on a network interface.	Mbps	An abnormally high rate of outgoing traffic may require additional analysis.
Maximum bandwidth:	An estimate of the capacity of a network interface.	Mbps	
Bandwidth usage:	Indicates the percentage of	Percent	By comparing the bandwidth usage with

Measurement	Description	Measurement Unit	Interpretation
	bandwidth used by a network interface.		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.

Note:

If the Windows Network Traffic - VM test is not reporting measures for a guest, make sure that you have enabled the SNMP service for the guest.

2.5.8 Network Traffic - VM Test

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

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every *Linux virtual_guest:network_interface* combination

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.

4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

   - **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Incoming network 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 network 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).

2.5.9 Tcp - VM Test

This test tracks various statistics pertaining to TCP connections to and from each guest of a XenServer host. The details of the test are provided below:

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each powered-on guest on the XenServer being monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed

2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.
8. **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*.

9. **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 **2.7.2**.

    - **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Incoming connections to VM:	Indicates the connections per second received by the guest.	Conns/Sec	A high value can indicate an increase in input load.
Outgoing connections to VM:	Indicates the connections per second initiated by the guest.	Conns/Sec	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.
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.

Measurement	Description	Measurement Unit	Interpretation
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.

2.5.10 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.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each powered-on guest on the XenServer being monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can

configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

- Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
- In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
- Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.
8. **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*.

9. **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 **2.7.2**.

   - **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 PASSWORD**s 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 **2.5.1.1**.

   - **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.

10. **SEGMENTS_SENT_MIN** - Specify the minimum threshold for the number of segments sent/transmitted over the network. The default value is 10; in this case, the test will compute/report the *Retransmit ratio from VM* measure only if more than 10 segments are sent over the network – i.e., if the value of the Segments sent by VM measure crosses the value 10. On the other hand, if the *Segments sent by VM* measure reports a value less than 10, then the test will not compute/report the *Retransmit ratio from VM* measure. This is done to ensure that no false alerts are generated by the eG Enterprise system for the *Retransmit ratio from VM* measure. You can change this minimum threshold to any value of your choice.

11. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

12. **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.

13. **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.

14. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

15. **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 *virtualmachine* 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.

16. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
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

Measurement	Description	Measurement Unit	Interpretation
			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.

2.5.11 Handles Usage - VM Test

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

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each powered-on guest on the Citrix XenServer monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER.** If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

    - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.

- In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
- Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.
8. **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*.
9. **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 **2.7.2**.
- **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 PASSWORD**s 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 **2.5.1.1**.
- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **HANDLES GROWTH LIMIT** - This defines the upper limit of the handles opened by any process. By default, this parameter is set to *8000*.

12. **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.

13. **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.

14. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

15. **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 *virtualmachine* 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.

16. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

17. **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**.

18. **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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
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.	Number	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.
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

Measurement	Description	Measurement Unit	Interpretation
			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.

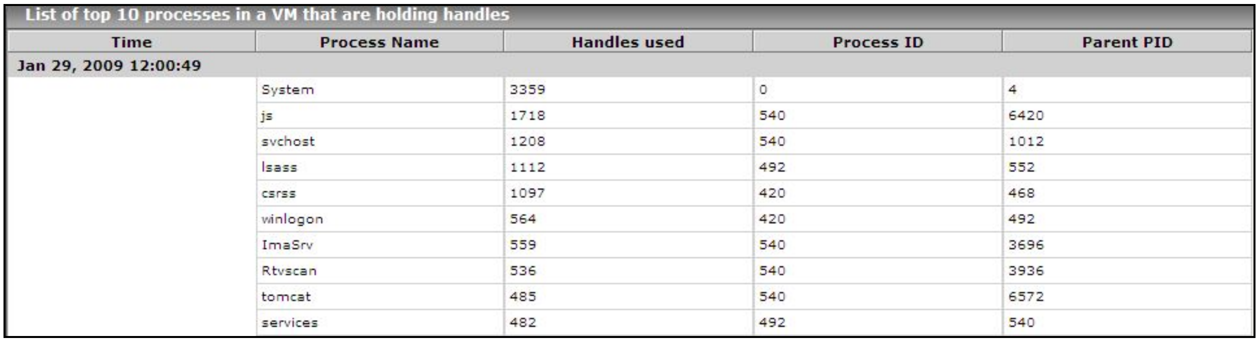

List of top 10 processes in a VM that are holding handles

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 2.30: 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.

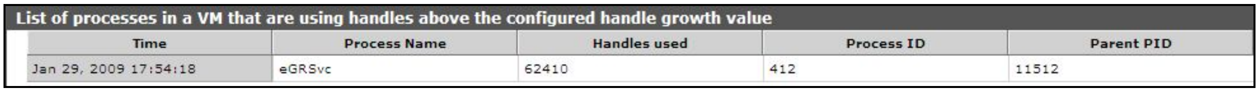

List of processes in a VM that are using handles above the configured handle growth value

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

Figure 2.31: The detailed diagnosis of the Processes using handles above limit in VM measure

2.5.12 Windows Services - VM Test

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

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every combination of *virtual_guest:disk_partition*

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.

3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.
8. **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*.

9. **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 **2.7.2**.

   - **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

16. **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.

17. **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**.

18. **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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
New automatic services started:	Indicates the number of Windows services with startup type as automatic, which were running in the	Number	The detailed diagnosis of this measure lists the services (with startup type as automcatic) that are running.

Measurement	Description	Measurement Unit	Interpretation
	last measurement period.		
New automatic services stopped:	Indicates the number of Windows services with startup type as automatic, which were not running in the last measurement period.	Number	To know which services stopped, use the detailed diagnosis of this measure (if enabled).
New manual services started:	Indicates the number of Windows services with startup type as manual, which were running in the last measurement period.	Number	Use the detailed diagnosis of this measure to identify the manual services that are running.
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.

2.5.13 GPU - VM Test

GPU-accelerated computing is the use of a graphics processing unit (GPU) together with a CPU to accelerate scientific, analytics, engineering, consumer, and enterprise applications. GPU- accelerated computing enhances application performance by offloading compute-intensive portions of the application to the GPU, while the remainder of the code still runs on the CPU. Architecturally, while a CPU has only few cores and handles few hundred threads at a time, a GPU is composed of hundreds of cores that can handle thousands of threads simultaneously and render a flawless rich graphics experience.

Now, imagine if you could access your GPU-accelerated applications, even those requiring intensive graphics power, anywhere on any device. **NVIDIA GRID** makes this possible. With NVIDIA GRID, a virtualized GPU designed specifically for virtualized server environments, data center managers can bring true PC graphics-rich experiences to users.

The NVIDIA GRID GPUs will be hosted in enterprise data centers and allow users to run virtual desktops or virtual applications on multiple devices connected to the internet and across multiple operating systems, including PCs, notebooks, tablets and even smartphones. Users can utilize their online-connected devices to enjoy the GPU power remotely.

In VDI/virtualized server environments, the NVIDIA GRID delivers GPU resources to virtual desktops/VMs. This way, graphics can be rendered on a virtual machine's (VM's) host server rather than on a physical end-point device. This technology now makes it possible to use virtual desktop technology to support users accessing graphics intensive workloads. There are two modes of making GPU resources available to virtual desktops:

- **Dedicated GPU or GPU Pass-through Technology:** NVIDIA GPU pass-through technology lets you create a virtual workstation that gives users all the benefits of a dedicated graphics processor at their desk. By directly connecting a dedicated GPU to a virtual machine through the hypervisor, you can now allocate the full GPU and graphics memory capability to a single virtual machine without any resource compromise.

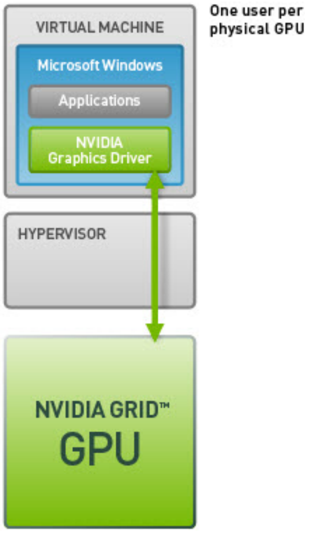

Figure 2.32: Dedicated GPU Technology

- **Shared GPU or Virtual GPU (vGPU) Technology:** GRID vGPU is the industry's most advanced technology for sharing true GPU hardware acceleration between multiple virtual desktops—without compromising the graphics experience. With GRID vGPU technology, the graphics commands of each virtual machine are passed directly to the GPU, without translation by the hypervisor. This allows the GPU hardware to be time-sliced to deliver improved shared virtualized graphics performance. The GRID vGPU manager allows for management of user profiles. IT managers can assign the optimal amount of graphics memory and deliver a customized graphics profile to meet the specific needs of each user. Every virtual desktop has dedicated graphics memory, just like they would at their desk, so they always have the resources they need to launch and run their applications.

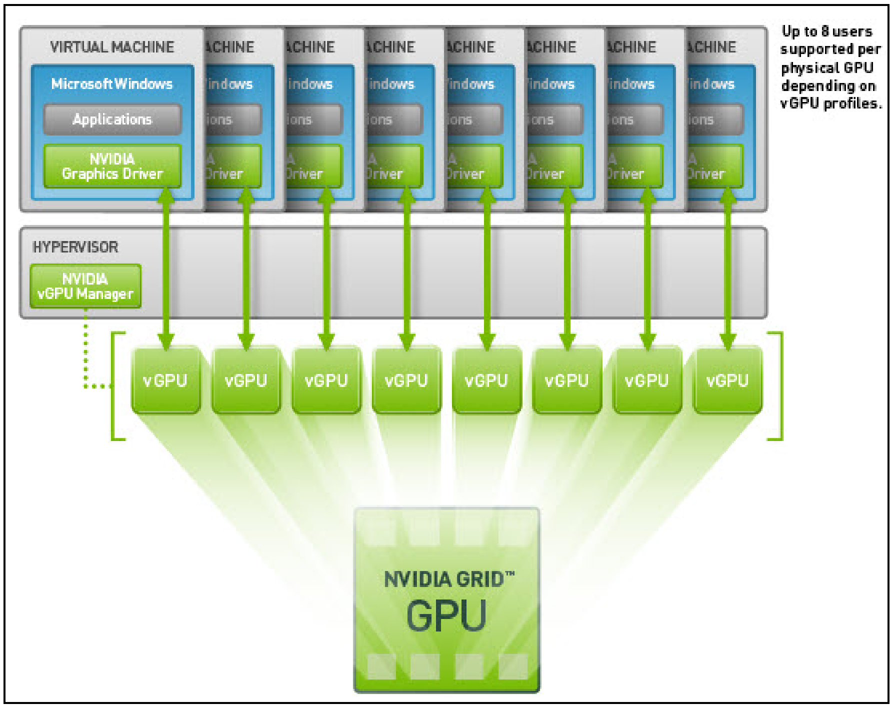

Figure 2.33: Shared vGPU Technology

In GPU-enabled Citrix XenServer environments, if users to VMs/virtual desktops complain of slowness when accessing graphic applications, administrators must be able to instantly detect the slowness and figure out its root-cause –is it because adequate GPU resources are not allocated to the VMs/virtual desktops? Is it because of excessive utilization of GPU memory and processing resources by a few VMs/virtual desktops? Or is it because the GPU clock frequencies are improperly set for one/more GPUs used by a VM/virtual desktop? Accurate answers to these questions can help administrators determine whether/not:

- The VMs/virtual desktops have been allocated enough vGPUs;
- The vGPUs are configured with enough graphics memory;
- The vGPU clock frequencies are rightly set;
- The GPU technology in use – i.e., the GPU Pass-through technology or the Shared vGPU technology – is ideal for the graphics processing requirements of the environment;

Measures to right-size the VMs/virtual desktops and fine-tune their GPU configuration can be initiated based on the results of this analysis. This is exactly what the **GPU – VM** test helps administrators achieve!

This test tracks the rate at which each vGPU processes frames, and thus pinpoints those vGPUs that are experiencing a processing bottleneck. The test also monitors the memory usage on each vGPU and helps administrators identify the vGPUs where memory is over-used. The test also reveals how each of these VMs/virtual desktops use each of the allocated vGPUs, thus enabling administrators to determine whether/not the allocated vGPUs are sufficient for the current and future processing requirements of the VMs/virtual desktops. In the process, the test also pinpoints those VMs/virtual desktops that are over-utilizing the graphical processors assigned to them. Also, to make sure that the assigned vGPUs are functioning without a glitch, the power consumption, temperature, and clock frequency of each vGPU is also checked at periodic intervals, so that abnormalities can be quickly detected.

Note:

This test will report metrics for only those Windows VMs where the **NVWMI** is installed. The steps for installing **NVWMI** and configuring the eG agent to use it have been detailed in Section 1.5 of this document.

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results every vGPU assigned to each Windows VM on the XenServer being monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

    - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
    - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
    - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
    - Then, click the **Update** button to save the changes.

    Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent

communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

- Obtain the server-certificate for the XenServer
- Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

16. **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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Cooler rate:	Indicates the percentage of device cooler rate for this GPU of this VM/virtual desktop.	Percentage	

Measurement	Description	Measurement Unit	Interpretation
Virtual GPU compute utilization:	Indicates the proportion of time over the past sample period during which one or more kernels were executing on this vGPU of this VM/virtual desktop.	Percentage	A value close to 100% indicates that the GPU of the VM/virtual desktop is busy processing graphic requests almost all the time. In a Shared vGPU environment a vGPU may be in use almost all the time, if the VM/virtual desktop it is allocated to run graphic- intensive applications. A resource-hungry VM/virtual desktop on a XenServer can impact the performance of other VMs/virtual desktops on the same server. If you find that only a single VM/virtual desktop has been consistently hogging the GPU resources, you may want to switch to the Dedicated GPU mode, so that excessive GPU usage by that VM/virtual desktop has no impact on the performance of other VMs/virtual desktops on that host. If all GPUs assigned to a VM/virtual desktop are found to be busy most of the time, you may want to consider allocating more GPU resources to that VM/virtual desktop.
Power consumption:	Indicates the current power usage of this GPU allocated to this VM/virtual desktop.	Watts	A very high value is indicative of excessive power usage by the GPU. Compare the value of this measure across GPUs to know which VM's/virtual desktop's GPU is consuming power excessively.
Core GPU temperature:	Indicates the current temperature of this GPU allocated to this VM/virtual desktop.	Celsius	Ideally, the value of this measure should be low. A very high value is indicative of abnormal GPU temperature. Compare the value of this measure across VMs/virtual desktops to identify that VM/virtual desktop for which GPU

Measurement	Description	Measurement Unit	Interpretation
			temperature soared since the last reading. To reduce the heat output of the GPU and consequently its temperature, you may consider performing underclocking. For instance, it is possible to set a GPU to run at lower clock rates when performing everyday tasks (e.g. internet browsing and word processing), thus allowing the card to operate at lower temperature and thus lower, quieter fan speeds.
Total framebuffer memory:	Indicates the total size of frame buffer memory of this GPU of this VM/virtual desktop.	MB	Frame buffer memory refers to the memory used to hold pixel properties such as color, alpha, depth, stencil, mask, etc.
Used frame buffer memory:	Indicates the amount of frame buffer memory on-board this GPU that has been used by this VM/virtual desktop.	MB	Properties like the screen resolution, color level, and refresh speed of the frame buffer can impact graphics performance. Also, if Error-correcting code (ECC) is enabled, the frame buffer memory usage will increase by several percent. This is because, ECC uses up memory to detect and correct the most common kinds of internal data corruption. Moreover, the driver may also reserve a small amount of memory for internal use, even without active work on the GPU; this too may impact frame buffer memory. For optimal graphics performance therefore, adequate frame buffer memory should be allocated to the VM/virtual desktop.

Measurement	Description	Measurement Unit	Interpretation
Free frame buffer memory:	Indicates the amount of frame buffer memory on-board this GPU that has not been yet been used by this VM/virtual desktop.	MB	
Frame buffer memory utilization:	Indicates the percentage of total frame buffer memory that has been allocated to this VM/virtual desktop.	Percentage	Ideally, the value of this measure should be low. A value close to 100% is indicative of excessive usage of frame buffer memory. Properties like the screen resolution, color level, and refresh speed of the frame buffer can impact graphics performance. Also, if Error-correcting code (ECC) is enabled, the frame buffer memory usage will increase by several percent. This is because, ECC uses up memory to detect and correct the most common kinds of internal data corruption. Moreover, the driver may also reserve a small amount of memory for internal use, even without active work on the GPU; this too may impact frame buffer memory usage. For optimal graphics performance therefore, adequate frame buffer memory should be allocated to the VM/virtual desktop.
Virtual memory:	Indicates the virtual memory of this GPU device of this VM/virtual desktop.	MB	
GPU memory utilization:	Indicates the percentage of time over the past sample	Percentage	A value close to 100% is a cause for concern as it indicates that the graphics

Measurement	Description	Measurement Unit	Interpretation
	period during which memory on this GPU was read/written on by this VM/virtual desktop.		memory on a GPU is almost always in use. In a Shared vGPU environment, memory may be consumed all the time if one/more VMs/virtual desktops utilize the graphics memory excessively and constantly. If you find that only a single VM/virtual desktop has been consistently hogging the graphic memory resources, you may want to switch to the Dedicated GPU mode, so that excessive memory usage by that VM/virtual desktop has no impact on the performance of other VMs/virtual desktops on that host. If the value of this measure is high almost all the time for most of the GPUs, it could mean that the VM/virtual desktop is not sized with adequate graphics memory.
Total BAR1 memory:	Indicates the total size of the BAR1 memory of this GPU allocated to this VM/virtual desktop.	MB	BAR1 is used to map the frame buffer (device memory) so that it can be directly accessed by the CPU or by 3rd party devices (peer-to-peer on the PCIe bus).
Used BAR1 memory:	Indicates the amount of BAR1 memory on this GPU that is used by this VM/virtual desktop.	MB	For better user experience with graphic applications, enough BAR1 memory should be available to the VM/virtual desktop.
Free BAR1 memory:	Indicates the total size of BAR1 memory of this GPU that is yet to be used by this VM/virtual desktop.	MB	
BAR1 memory utilization:	Indicates the percentage of the allocated BAR1 memory that is currently	Percentage	A value close to 100% is indicative of excessive usage of the BAR1 memory by a VM/virtual desktop.

Measurement	Description	Measurement Unit	Interpretation
	being utilized by this VM/virtual desktop.		For best graphics performance, this value should be low. To ensure that, adequate BAR1 memory should be allocated to the VM.
Power management:	Indicates whether/not power management is enabled for this GPU of this VM/virtual desktop.		Many NVIDIA graphics cards support multiple performance levels so that the server can save power when full graphics performance is not required. The default Power Management Mode of the graphics card is Adaptive. In this mode, the graphics card monitors GPU usage and seamlessly switches between modes based on the performance demands of the application. This allows the GPU to always use the minimum amount of power required to run a given application. This mode is recommended by NVIDIA for best overall balance of power and performance. If the power management mode is set to Adaptive, the value of this measure will be Supported. Alternatively, you can set the Power Management Mode to Maximum Performance. This mode allows users to maintain the card at its maximum performance level when 3D applications are running regardless of GPU usage. If the power management mode of a GPU is Maximum Performance, then the value of this measure will be Maximum. The numeric values that correspond to these measure values are discussed in the table below:

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td><table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Supported</td><td>1</td></tr><tr><td>Maximum</td><td>0</td></tr></table><b>Note:</b>
By default, this measure will report the <b>Measure Value</b>s listed in the table above to indicate the power management status. In the graph of this measure however, the same is represented using the numeric equivalents only.</td></tr>
<tr><td><b>Power limit:</b></td><td>Indicates the power limit configured for this GPU of this VM/virtual desktop.</td><td>Watts</td><td><b>This measure will report a value only if the value of the ‘Power management’ measure is ‘Supported’.</b>
The power limit setting controls how much voltage a GPU can use when under load. Its not advisable to set the power limit at its maximum – i.e., the value of this measure should not be the same as the value of the Max power limit measure - as it can cause the GPU to behave strangely under duress.</td></tr>
<tr><td><b>Default power limit:</b></td><td>Indicates the default power management algorithm's power ceiling for this GPU.</td><td>Watts</td><td><b>This measure will report a value only if the value of the ‘Power management’ measure is ‘Supported’.</b></td></tr>
<tr><td><b>Enforced power limit:</b></td><td>Indicates the power management algorithm's power ceiling for this GPU of this VM/virtual desktop.</td><td>Watts</td><td><b>This measure will report a value only if the value of the ‘Power management’ measure is ‘Supported’.</b>
The total board power draw is manipulated by the power management algorithm such that it stays under the value reported by this measure.</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
Min power limit:	The minimum value that the power limit be set to for this GPU of this VM/virtual desktop.	Watts	**This measure will report a value only if the value of the 'Power management' measure is 'Supported'.**
Max power limit:	The maximum value that the power limit for this GPU of this VM/virtual desktop can be set to.	Watts	**This measure will report a value only if the value of the 'Power management' measure is 'Supported'.** If the value of this measure is the same as that of the Power limit measure, then the GPU may behave strangely.
Core clock:	Indicates current frequency of the graphics clock on this GPU of this VM/virtual desktop.	MHz	GPU has many more cores than your average CPU but these cores are much simpler and much smaller so that many more actually fit on a small piece of silicon. These smaller, simpler cores go by different names depending upon the tasks they perform. Stream processors are the cores that perform a single thread at a slow rate. But since GPUs contain numerous stream processors, they make overall computation high. The streaming multiprocessor clock is how fast the stream processors run. The memory clock is how fast the memory on the card runs. The GPU core clock is the speed at which the GPU assigned to the VM/virtual desktop operates. By correlating the frequencies of these clocks – i.e., the value of these measures - with the memory usage, power usage, and overall performance of the GPU, you can figure out if overclocking is required or not.

Measurement	Description	Measurement Unit	Interpretation
			Overclocking is the process of forcing a GPU core/memory to run faster than its manufactured frequency. Overclocking can have both positive and negative effects on GPU performance. For instance, memory overclocking helps on cards with low memory bandwidth, and with games with a lot of post-processing/textures/filters like AA that are VRAM intensive. On the other hand, speeding up the operation frequency of a shader/streaming processor/memory clock, without properly analyzing its need and its effects, may increase its thermal output in a linear fashion. At the same time, boosting voltages will cause the generated heat to sky rocket. If improperly managed, these increases in temperature can cause permanent physical damage to the core/memory or even “heat death”. Putting an adequate cooling system into place, adjusting the power provided to the GPU, monitoring your results with the right tools and doing the necessary research are all critical steps on the path to safe and successful overclocking.
Memory clock:	Indicates current memory clock frequency on this GPU of this VM/virtual desktop.	MHz	
Streaming multiprocessor clock:	Indicates the current frequency of the streaming multiprocessor clock on this GPU of this VM/virtual desktop.	MHz	
Frame rate:	Indicates the rate at which frames are processed by this GPU of this VM/virtual desktop.	Frames/Sec	FPS is how fast your graphics card can output individual frames each second. It is the most time- tested and ideal measure of performance of a GPU. Higher the value of this measure, healthier is the GPU.
Fan speed:	Indicates the percent of maximum speed that this GPU’s fan is currently intended to run at.	Percentage	The value of this measure could range from 0 to 100%. An abnormally high value for this measure could indicate a problem

Measurement	Description	Measurement Unit	Interpretation
			condition – eg., a sudden surge in the temperature of the GPU that could cause the fan to spin faster. Note that the reported speed is only the intended fan speed. If the fan is physically blocked and unable to spin, this output will not match the actual fan speed. Many parts do not report fan speeds because they rely on cooling via fans in the surrounding enclosure. By default the fan speed is increased or decreased automatically in response to changes in temperature.
Compute processes:	Indicates the number of processes having compute context on this GPU of this VM.	Number	Use the detailed diagnosis of this measure to know which processes are currently using the GPU. The process details provided as part of the detailed diagnosis include, the PID of the process, the process name, and the GPU memory used by the process. Note that the GPU memory usage of the processes will not be available in the detailed diagnosis, if the Windows platform on which XenApp operates is running in the WDDM mode. In this mode, the Windows KMD manages all the memory, and not the NVIDIA driver. Therefore, the NVIDIA SMI commands that the test uses to collect metrics will not be able to capture the GPU memory usage of the processes.
Volatile single bit errors:	Indicates the number of volatile single bit errors in this GPU of this VM/virtual desktop.	Number	Volatile error counters track the number of errors detected since the last driver load. Single bit ECC errors are automatically corrected by the hardware and do not result in data corruption. Ideally, the value of this measure should

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>be 0.</td></tr>
<tr><td><b>Volatile double bit errors:</b></td><td>Indicates the total number of volatile double bit errors in this GPU of this VM/virtual desktop.</td><td>Number</td><td>Volatile error counters track the number of errors detected since the last driver load. Double bit errors are detected but not corrected.
Ideally, the value of this measure should be 0.</td></tr>
<tr><td><b>Aggregate single bit errors:</b></td><td>Indicates the total number of aggregate single bit errors in this GPU of this VM/virtual desktop.</td><td>Number</td><td>Aggregate error counts persist indefinitely and thus act as a lifetime counter. Single bit ECC errors are automatically corrected by the hardware and do not result in data corruption.
Ideally, the value of this measure should be 0.</td></tr>
<tr><td><b>Aggregate double bit errors:</b></td><td>Indicates the total number of aggregate double bit errors in this GPU of this VM/virtual desktop.</td><td>Number</td><td>Aggregate error counts persist indefinitely and thus act as a lifetime counter. Double bit errors are detected but not corrected.
Ideally, the value of this measure should be 0.</td></tr>
<tr><td><b>Mode:</b></td><td>Indicates the mode using which the GPU resources were delivered to this VMs.</td><td></td><td>The values that this measure can take and their corresponding numeric values are as follows:
<table>
<tr><th>Measure Value</th><th>Numeric Values</th></tr>
<tr><td>Pass through</td><td>0</td></tr>
<tr><td>Shared</td><td>1</td></tr>
<tr><td>Unavailable (GPU card is not allocated to any VM)</td><td>2</td></tr>
</table>
<b>Note:</b></td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			By default, this test reports the **Measure Value** s listed in the table above to indicate the mode of GPU delivery. In the graph of this measure however, the same is represented using the numeric equivalents only.
Physical GPU compute utilization:	Indicates the proportion of time over the past sample period during which one or more kernels were executing on the physical GPU of this VM/virtual desktop.	Percentage	**This measure will report metrics only VMs configured with a Tesla GPU card.** A value close to 100% indicates that the physical GPU is busy processing graphic requests from this VM almost all the time. In a Shared vGPU environment a vGPU may be in use almost all the time, if the VM/virtual desktop it is allocated to run graphic- intensive applications. A resource-hungry VM/virtual desktop on a XenServer can impact the performance of other VMs/virtual desktops on the same server. If you find that only a single VM/virtual desktop has been consistently hogging the GPU resources, you may want to switch to the Dedicated GPU mode, so that excessive GPU usage by that VM/virtual desktop has no impact on the performance of other VMs/virtual desktops on that host. If all GPUs assigned to a VM/virtual desktop are found to be busy most of the time, you may want to consider allocating more GPU resources to that VM/virtual desktop.

As stated earlier, by default, clicking on the **Virtual Servers** layer of a managed *Citrix XenServer*, 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 **Virtual Servers** 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 **Virtual Servers** layer is clicked, the list of tests mapped to that layer appears, as depicted by Figure 2.34.

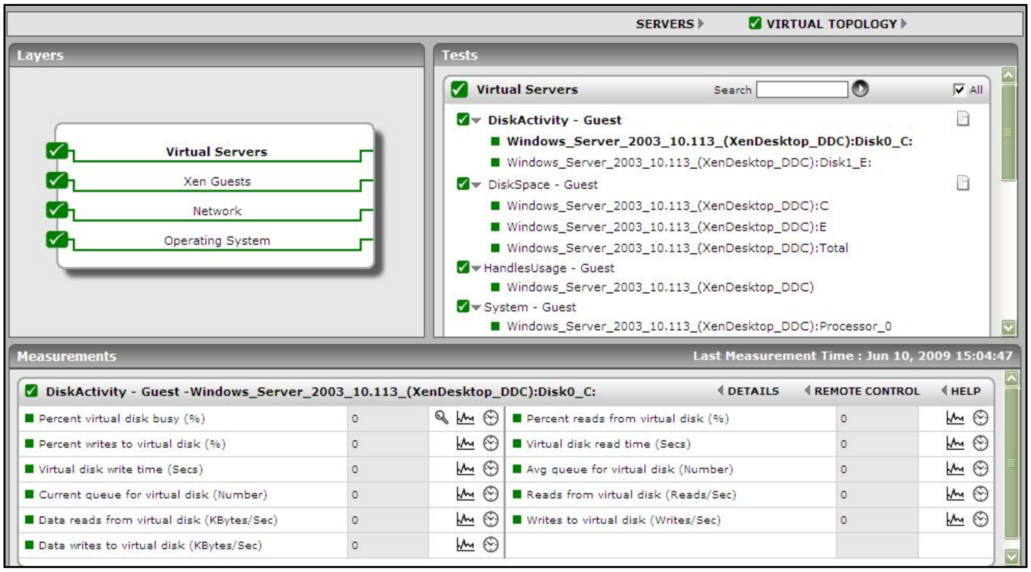

Virtual Servers

- DiskActivity - Guest
  - Windows_Server_2003_10.113_(XenDesktop_DDC):Disk0_C:
  - Windows_Server_2003_10.113_(XenDesktop_DDC):Disk1_E:
- DiskSpace - Guest
  - Windows_Server_2003_10.113_(XenDesktop_DDC):C
  - Windows_Server_2003_10.113_(XenDesktop_DDC):E
  - Windows_Server_2003_10.113_(XenDesktop_DDC):Total
- HandlesUsage - Guest
  - Windows_Server_2003_10.113_(XenDesktop_DDC)
- System - Guest
  - Windows_Server_2003_10.113_(XenDesktop_DDC):Processor_0

Last Measurement Time : Jun 10, 2009 15:04:47

DiskActivity - Guest -Windows_Server_2003_10.113_(XenDesktop_DDC):Disk0_C:

Percent virtual disk busy (%)	0	Percent reads from virtual disk (%)	0
Percent writes to virtual disk (%)	0	Virtual disk read time (Secs)	0
Virtual disk write time (Secs)	0	Avg queue for virtual disk (Number)	0
Current queue for virtual disk (Number)	0	Reads from virtual disk (Reads/Sec)	0
Data reads from virtual disk (KBytes/Sec)	0	Writes to virtual disk (Writes/Sec)	0
Data writes to virtual disk (KBytes/Sec)	0		

Figure 2.34: The tests mapped to the Virtual Servers layer

If you now want the **Server view** of Figure 2.20, simply click on the **SERVERS** link above the list of tests in Figure 2.34 (indicated by the arrow).

Clicking on any of the guests in the **Server view** leads you to Figure 2.35 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 2.35.

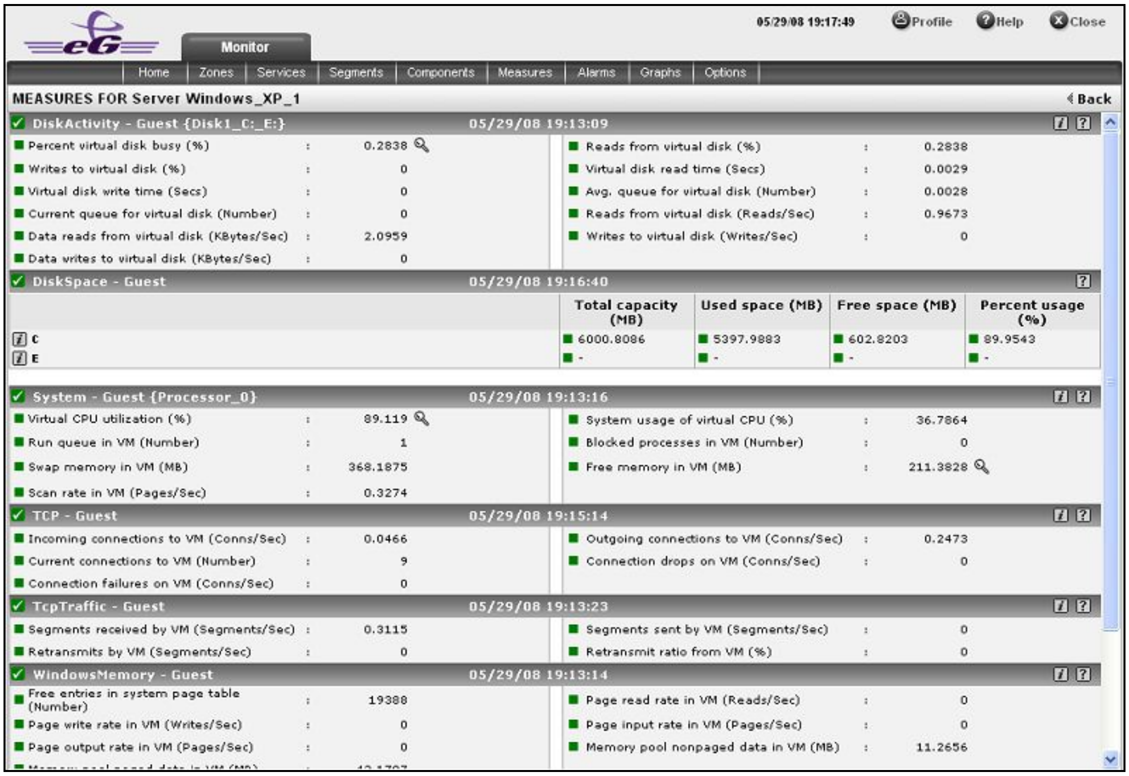

MEASURES FOR Server Windows_XP_1

✓ DiskActivity - Guest {Disk1_C:_E:} 05/29/08 19:13:09

Percent virtual disk busy (%)	:	0.2838	Reads from virtual disk (%)	:	0.2838
Writes to virtual disk (%)	:	0	Virtual disk read time (Secs)	:	0.0029
Virtual disk write time (Secs)	:	0	Avg. queue for virtual disk (Number)	:	0.0028
Current queue for virtual disk (Number)	:	0	Reads from virtual disk (Reads/Sec)	:	0.9673
Data reads from virtual disk (KBytes/Sec)	:	2.0959	Writes to virtual disk (Writes/Sec)	:	0
Data writes to virtual disk (KBytes/Sec)	:	0			

✓ DiskSpace - Guest 05/29/08 19:16:40

	Total capacity (MB)	Used space (MB)	Free space (MB)	Percent usage (%)
C	6000.8086	5397.9883	602.8203	89.9543
E	-	-	-	-

✓ System - Guest {Processor_0} 05/29/08 19:13:16

Virtual CPU utilization (%)	:	89.119	System usage of virtual CPU (%)	:	36.7864
Run queue in VM (Number)	:	1	Blocked processes in VM (Number)	:	0
Swap memory in VM (MB)	:	368.1875	Free memory in VM (MB)	:	211.3828
Scan rate in VM (Pages/Sec)	:	0.3274			

✓ TCP - Guest 05/29/08 19:15:14

Incoming connections to VM (Conns/Sec)	:	0.0466	Outgoing connections to VM (Conns/Sec)	:	0.2473
Current connections to VM (Number)	:	9	Connection drops on VM (Conns/Sec)	:	0
Connection failures on VM (Conns/Sec)	:	0			

✓ TcpTraffic - Guest 05/29/08 19:13:23

Segments received by VM (Segments/Sec)	:	0.3115	Segments sent by VM (Segments/Sec)	:	0
Retransmits by VM (Segments/Sec)	:	0	Retransmit ratio from VM (%)	:	0

✓ WindowsMemory - Guest 05/29/08 19:13:14

Free entries in system page table (Number)	:	19388	Page read rate in VM (Reads/Sec)	:	0
Page write rate in VM (Writes/Sec)	:	0	Page input rate in VM (Pages/Sec)	:	0
Page output rate in VM (Pages/Sec)	:	0	Memory pool nonpaged data in VM (MB)	:	11.2656

Figure 2.35: Measures pertaining to a chosen guest

To view real-time graphs of pre-configured measures (pertaining to the Xen host and the guests operating on it), click on the **LIVE GRAPH** link in Figure 2.20. Alternatively, you can click on the icon that appears in the **Tests** panel of the layer model page when the **Virtual Servers** layer is clicked to view the live graph. The graph display that appears subsequently (see Figure 2.36) has been organized in such a way that next to every host-pertinent measure graph, the closely related guest-specific measure graph appears. For instance, next to the graph of the '*Physical CPU usage*' measure of the *XenCpu* test, you will find a graph of the '*Physical CPU used*' measure of the *XenGuestDetails* test. This way, you can easily compare and correlate how well the physical CPU resources are being utilized by both the Xen host 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 XenServer host? or is it the virtual guest? If you access this page from the **LIVE GRAPH** link in Figure 2.20, then, by default, you will view live graphs pertaining to the *Citrix XenServer*. However, you can select a different virtualized component-type and a different virtualized component using the **Type** and **Component Name** lists (respectively) in Figure 2.36.

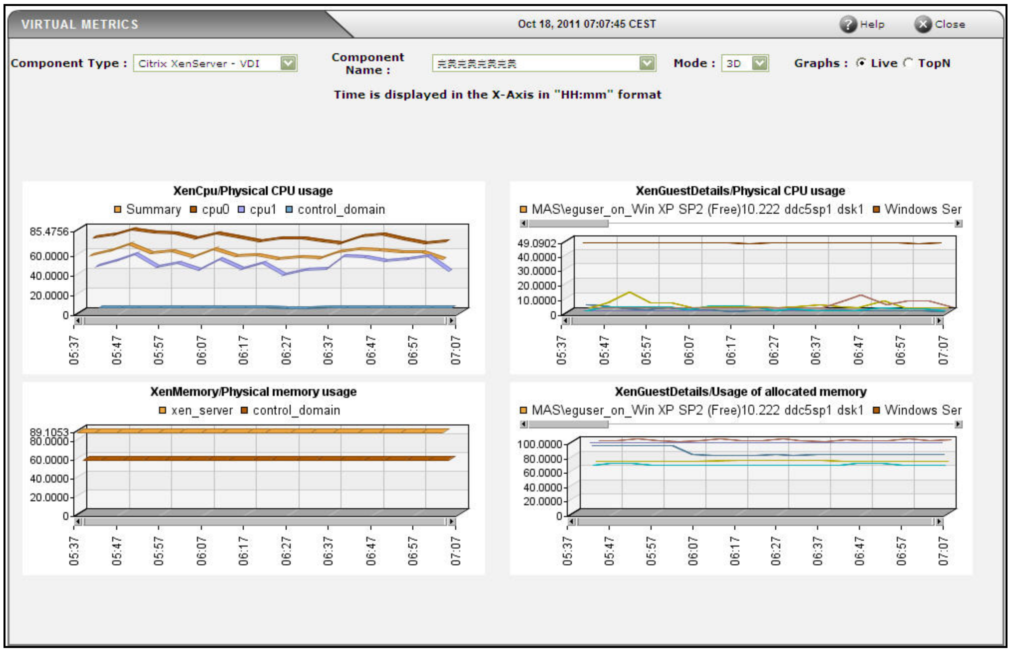

VIRTUAL METRICS

Oct 18, 2011 07:07:45 CEST

Component Type : Citrix XenServer - VDI

Component Name : [illegible]

Mode : 3D

Graphs : ● Live ○ TopN

Time is displayed in the X-Axis in "HH:mm" format

Figure 2.36: Live graph of the XenServer

2.5.14 Crash Details - VM Test

Event logs on Windows VMs capture critical error conditions such as service crashes and application crashes on the VMs, application and service hangs, and service errors. Since the crash/slowness experienced by any mission-critical program/service on a Windows VM may affect the uptime of the dependent business services, administrators should be able to instantly capture these serious problem conditions, investigate the reasons for their occurrence, and promptly resolve them. This is exactly what the Crash Details -VM test helps administrators achieve! This test periodically scans the event logs on each Windows VM and reports the count of crashes, hangs, and errors that may have occurred recently on that VM. Detailed diagnostics provided by this test pinpoints the applications/services that crashed, hanged, or encountered errors, and thus enables quick and efficient troubleshooting.

Note:

This test will not report metrics on VMs running Windows 2000/2003/XP.

Target of the test : A Citrix XenServer

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every Windows VM on the Citrix XenServer being monitored.

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.

3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.

4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH.**
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.

6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.

7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

   - **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

14. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using

the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

15. **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 *virtualmachine* 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.

16. **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**.

17. **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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Recent application crashes	Indicates the number of application crash events that occurred on this VM during the last measurement period.	Number	An event with the ID 1000 is logged in the event log every time a program terminates unexpectedly on a virtual desktop. This measure reports the number of events in the event log with event ID 1000. Use the detailed diagnosis of this measure to know which programs and modules stopped suddenly.

Measurement	Description	Measurement Unit	Interpretation
Recent service crashes	Indicates the number of service crash events that occurred on thisVM during the last measurement period.	Number	An event with the ID 7031 is logged in the Service Control Manager every time a service terminates ungracefully. This measure reports the number of events in the event log with event ID 7031. Use the detailed diagnosis of this measure to know the complete details of such events.
Recent application hangs	Indicates the number of application hang events that occurred on this VM during the last measurement period.	Number	An event with the ID 1002 is logged in the Application Event Log every time an application hangs. This measure reports the number of events in the event log with event ID 1002. Use the detailed diagnosis of this measure to know the complete details of the recent application hang events.
Recent service hangs	Indicates the number of service hang events that occurred on this VM during the last measurement period.	Number	An event with the ID 7022 is logged in the Service Control Manager every time a service hangs. This measure reports the number of events in the event log with event ID 7022. Use the detailed diagnosis of this measure to know the complete details of the recent service hang events.
Recent service errors	Indicates the number of service errors that occurred on this VM during the last measurement period.	Number	Events with the ID 7023, 7024, and 7026 are logged in the Service Control Manager every time a service error occurs. This measure reports the number of events in the event log with the aforesaid event IDs. Use the detailed diagnosis of this measure to know the complete details of the recent service errors.

2.5.15

2.6 Correlation Between Applications in a Xen Virtualized Environments

Using the eG Enterprise administration console, administrators can add applications running on the VMs for monitoring. To monitor these applications, agents can be installed in the guests, or an agentless monitoring approach can be used. To effectively monitor the applications running in a virtual environment, it is important to be able to determine which XenServer an application is running on. This mapping of applications to virtual servers is important for root-cause diagnosis - for example, a problem with the virtual server (e.g., excessive disk slowdowns) can impact the performance of all the applications running on the server’s virtual machines.

eG Enterprise is able to automatically determine the mapping of applications to virtual servers. Whether eG Enterprise automatically determines the mapping of applications to virtual servers or not is determined by the value of the **AutoVirtualMapping** variable in the **[MISC]** section of the **eg_external.ini** configuration file in the **<EG_INSTALL_DIR>\manager\config** directory of the eG manager. If the value of this variable is **true**, the eG manager auto-discovers the applications to virtual servers mapping.

Note:

- For **AutoVirtualMapping** to work, the detailed diagnosis frequencies set globally (i.e., using the Configure -> Diagnosis menu sequence) should not be set to *0:0*.
- As long as the **Identify agents only using nick names** flag in the **MANAGER SETTINGS** page of the eG administrative interface (Configure -> Manager Settings menu sequence) is **Yes** (which is the default), eG Enterprise can automatically identify the server applications executing on an XenServer host, using the host/nick names that are mapped to the IP addresses discovered on the host. If the **Identify agents only using nick names** flag is set to **No** instead, then make sure that, while managing a server application executing in a virtualized environment, the hostname of the virtual machine is specified as the nick name of the corresponding server application. If more than one server application is executing on the same virtual machine, then any one of those server applications should have the virtual machine name as its nick name.

To disable auto-discovery, set this value to **false**. In such a case, once a *Citrix XenServer* is added, then, when adding any new server application using the eG administrative interface, you will be prompted to manually set an association between the server application being added and the Virtual Server.

The mapping of applications to virtual servers is used by eG Enterprise for correlation - e.g., since the application runs on the virtual server, it is most likely that a problem with the virtual server will impact the performance of the application running on one of the guests. To view this application-virtual server association, simply click on the **VIRTUAL TOPOLOGY** link in the layer model page of the virtual server.

Note:

The **VIRTUAL TOPOLOGY** link will also be available in the layer model page of those server applications that are executing on virtual guests.

Doing so reveals Figure 2.37 depicting the *Citrix XenServer* and the server applications executing on it. By clicking on any of the components in Figure 2.37, the user can drill down into specific layers of this component for specific details on the performance of the component.

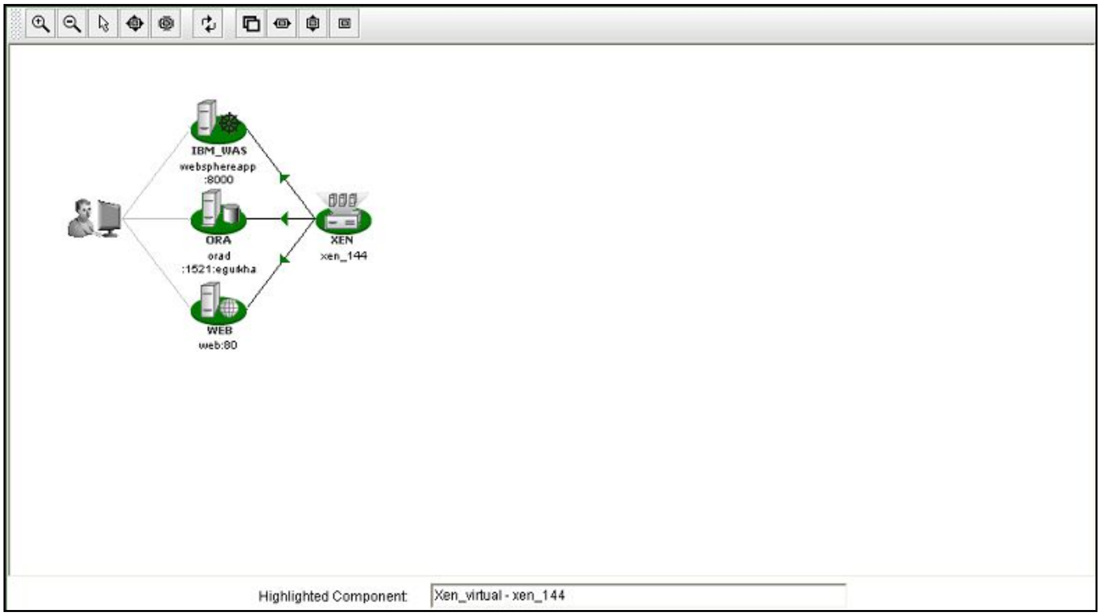

Highlighted Component: Xen_virtual - xen_144

Figure 2.37: Depicts the applications that have been deployed on the guest OS of a virtual server

The arrows in Figure 2.37 depict the dependencies between the virtual server host and the applications running on it. Since the applications are hosted on one of the guests running on the host, they depend on the virtual server host - i.e., any unusual resource usage on the virtual server host impacts the applications running on any of the virtual guests. The dependency information between the virtual server host and the applications hosted on it is used by eG Enterprise for end-to-end correlation.

2.7 Troubleshooting

2.7.1 Troubleshooting the Failure of the eG Agent to Report Measures for a XenServer

If the eG agent is not reporting any measures for a XenServer, then do the following:

1. Check the version of the XenServer being monitored; if the target XenServer is of version 5 (or above), then, do the following:

   - *Edit the* **eg_tests.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory).
   - In the **[AGENT_SETTINGS]** section of the file, the **XenPerfFromXml** flag is set to **Yes** by default. This implies that, by default, the eG agent collects metrics from the XenServer by establishing an HTTP/HTTPS connection with the server. Sometimes, the eG agent may be unable to access the XenServer via HTTP/HTTPS - say, owing to access restrictions imposed by a firewall configuration.

In such cases, set this flag to **No**, so that the eG agent uses the XenServer API instead to collect the required metrics.

- Then, login to the XenServer as root user.
- Run the following command on the console of the XenServer:

  xe host-list

  The above command will print the host details of the XenServer. Look for the UUID value. This is required in the next step. If your XenServer is part of a pool, the above command will list the details for all the servers in the pool. Look for the details of the current XenServer and find its UUID value.

- Next, run the following command on the console of the XenServer by specifying the UUID determined in the previous step:

  xe host-param-set uuid=<UUID acquired in the step above> other-config:rrd_update_interval=2

  Note that the **rrd_update_interval** can hold any of the below-mentioned values:

  - **never**-this is default, meaning no periodic polling is performed
  - **1** - Polling is performed every 5 seconds.
  - **2** - Polling is performed every minute (**we recommend this value for monitoring purposes**).

- Then, restart the XenAPI by issuing the following command from the **/opt/xensource/bin** directory(on Linux; on Windows, this will be: **<XEN_INSTALL_DIR>\xensource\bindirectory)**:

  xe-toolstack-restart

- Now, restart the eG agent and verify whether it is able to collect the metrics.

2. If the aforesaid steps do not produce the desired results, then check whether the XenServer has been configured to use SSL. If so, switch on the **SSL** flag of all tests to **Yes**.
3. If the tests still do not work, then, check whether the XenServer is using its default SSL certificate or has been configured to use a self-signed certificate. If a self-signed certificate is in use, then follow the broad steps below to make sure the eG agent communicates with the XenServer using this cerficate:

   a. Obtain the server-certificate for the XenServer

   b. Import the server-certificate into the local certificate store of the eG agent

The sections to come will discuss each of these steps in detail.

2.7.1.1 Obtaining the Server-Certificate for the XenServer

To obtain a certificate, you can login to the XenServer as any user with access privileges to the server, and copy the certificate directly from the server to the agent host.

If the agent is deployed on the control domain of the server, then follow the steps detailed below to obtain the server-certificate:

1. Login to the XenServer as a user with access rights to the server. If you have created a special user, then the credentials of that user can be supplied while logging into the server.
2. Since the certificate needs to be stored on the agent host, create a directory on the XenServer itself for storing the certificate.
3. Next, copy the certificate file (with the extension, **.pem**) from the **/etc/xensource** directory on the server to the directory that you created at step 2, using a unique filename (preferably, a name that suggests that the file contains an SSL certificate - say, *sslcert*).

If you intend to monitor the virtualized environment in an 'agentless' manner, then the eG agent will be deployed on a remote Windows host. In such a case, you first need to copy the certificate from the XenServer to the remote agent's host. In such a case, follow the steps detailed below to obtain the certificate from the XenServer:

1. Open the Internet Explorer on the remote agent's host.
2. Connect to the XenServer using the HTTPS protocol: **https:// <Xen_server_IP>**.
3. A **Security Alert** message displays a warning regarding the certificate's certifying authority (see Figure 2.38).

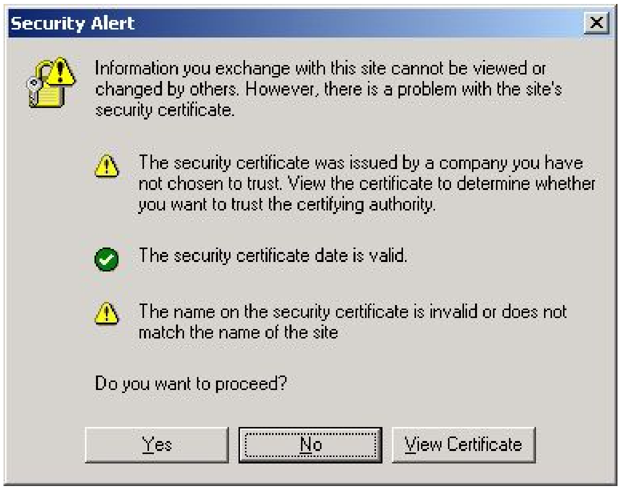

Security Alert

Information you exchange with this site cannot be viewed or changed by others. However, there is a problem with the site's security certificate.

The security certificate was issued by a company you have not chosen to trust. View the certificate to determine whether you want to trust the certifying authority.

The security certificate date is valid.

The name on the security certificate is invalid or does not match the name of the site

Do you want to proceed?

Figure 2.38: A Security Alert Message

4. Click on **View Certificate** in Figure 2.38 to open the **Certificate** properties page (see Figure 2.39).

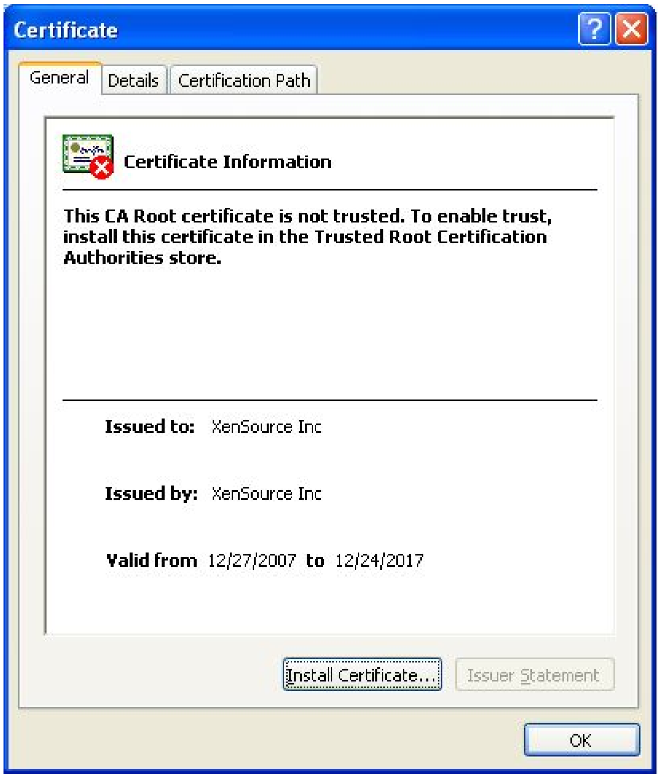

Certificate Information

This CA Root certificate is not trusted. To enable trust, install this certificate in the Trusted Root Certification Authorities store.

Issued to: XenSource Inc

Issued by: XenSource Inc

Valid from 12/27/2007 to 12/24/2017

Figure 2.39: Certificate properties page

5. Click **Install Certificate** in Figure 2.39 to launch the **Certificate Import Wizard** (see Figure 2.40).

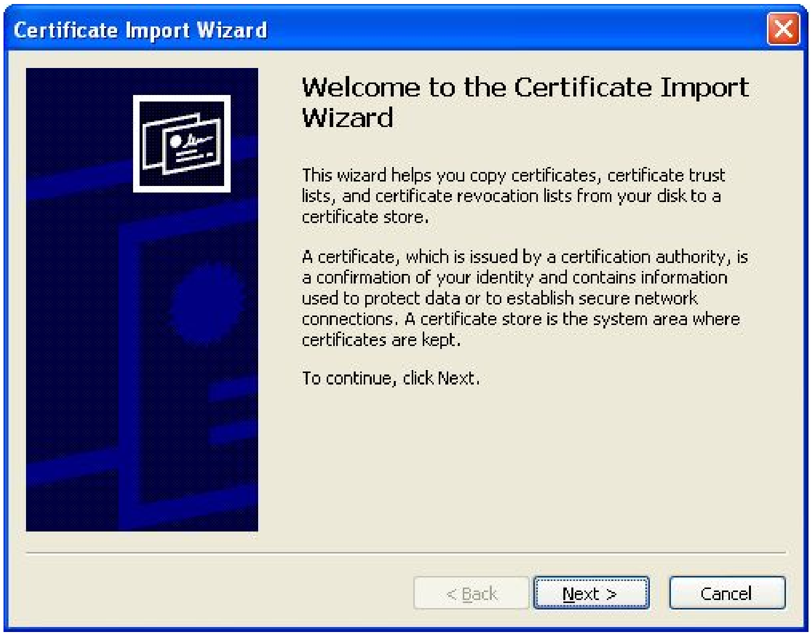

Welcome to the Certificate Import Wizard

This wizard helps you copy certificates, certificate trust lists, and certificate revocation lists from your disk to a certificate store.

A certificate, which is issued by a certification authority, is a confirmation of your identity and contains information used to protect data or to establish secure network connections. A certificate store is the system area where certificates are kept.

To continue, click Next.

Figure 2.40: The Certificate Import Wizard’s Welcome screen

6. In the **Welcome** screen, click **Next** to proceed (see Figure 2.40). In the next step (see Figure 2.41), keep

the default setting, **Automatically select the certificate store based on the type of certificate**, and click **Next** to continue.

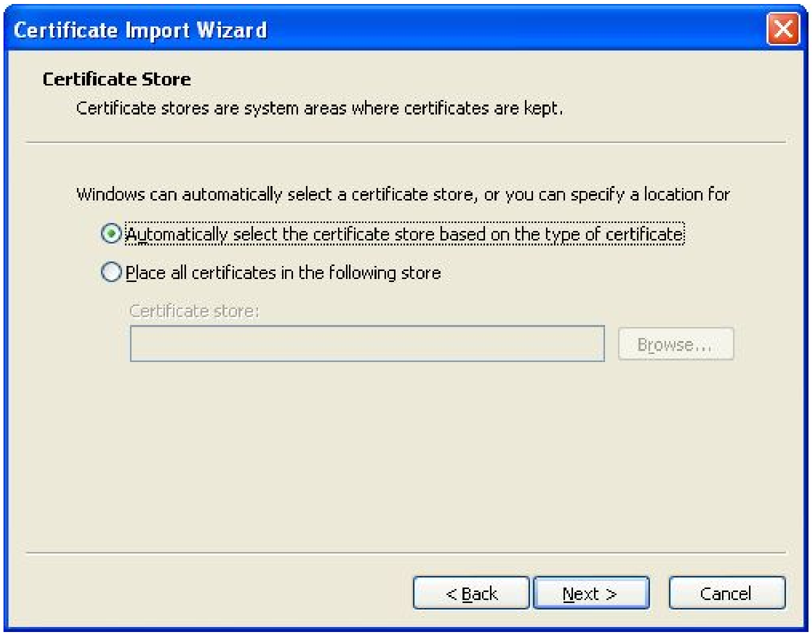

Certificate Import Wizard

Certificate Store

Certificate stores are system areas where certificates are kept.

Windows can automatically select a certificate store, or you can specify a location for

● Automatically select the certificate store based on the type of certificate

○ Place all certificates in the following store

Certificate store:

Figure 2.41: Specifying the certificate store

7. Next, click **Finish** to complete the import (see Figure 2.42).

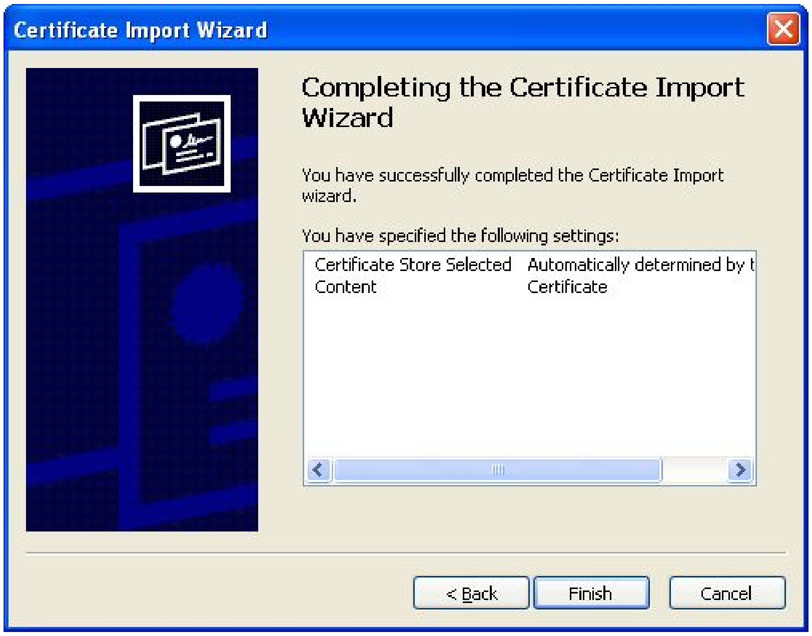

Certificate Import Wizard

Completing the Certificate Import Wizard

You have successfully completed the Certificate Import wizard.

You have specified the following settings:

Certificate Store Selected	Automatically determined by t
Content	Certificate

Figure 2.42: Finishing the certificate import

8. A **Root Certificate Store** alert message then appears, providing a summary of the details of the certificate that you imported (see Figure 2.43).

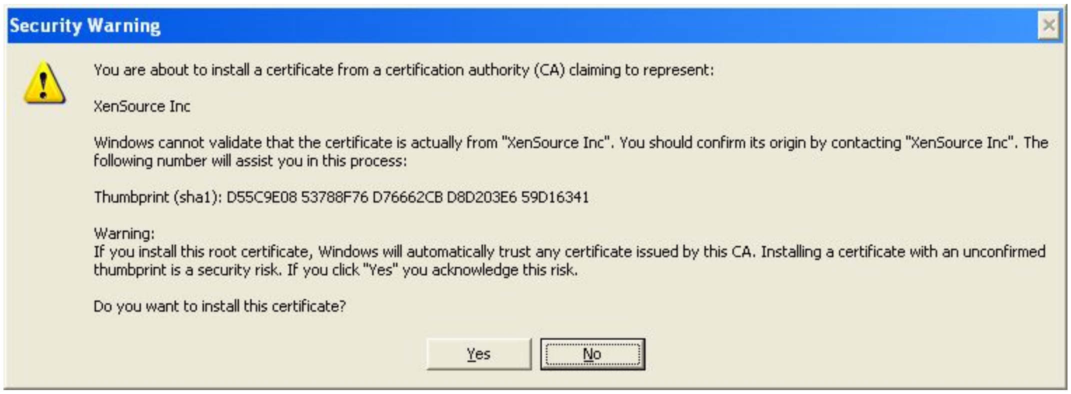

Security Warning

You are about to install a certificate from a certification authority (CA) claiming to represent:

XenSource Inc

Windows cannot validate that the certificate is actually from "XenSource Inc". You should confirm its origin by contacting "XenSource Inc". The following number will assist you in this process:

Thumbprint (sha1): D55C9E08 53788F76 D76662CB D8D203E6 59D16341

Warning:
If you install this root certificate, Windows will automatically trust any certificate issued by this CA. Installing a certificate with an unconfirmed thumbprint is a security risk. If you click "Yes" you acknowledge this risk.

Do you want to install this certificate?

Figure 2.43: Summary of the certificate

9. Click **Yes** in Figure 2.43 to continue with the certificate installation. Upon a successful import, the following message appears (see Figure 2.44).

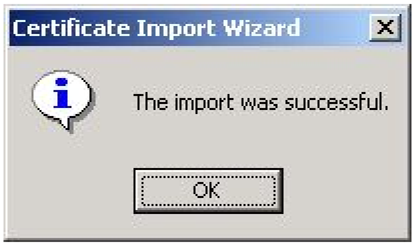

Certificate Import Wizard

The import was successful.

Figure 2.44: Successful import

10. Click **OK** to dismiss the success message (see Figure 2.44). The **Certificate** properties page becomes active again (see Figure 2.45).

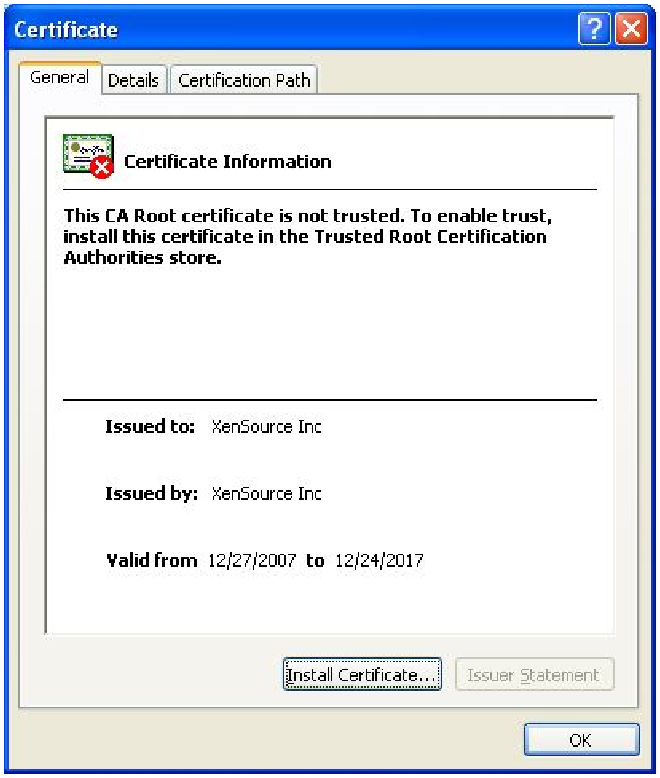

Certificate

Certificate Information

This CA Root certificate is not trusted. To enable trust, install this certificate in the Trusted Root Certification Authorities store.

Issued to: XenSource Inc

Issued by: XenSource Inc

Valid from 12/27/2007 to 12/24/2017

Figure 2.45: The Certificate properties page returns

11. Click **OK** in the **Certificate** dialog box of Figure 2.45 to continue. The **Security Alert Message** reappears.

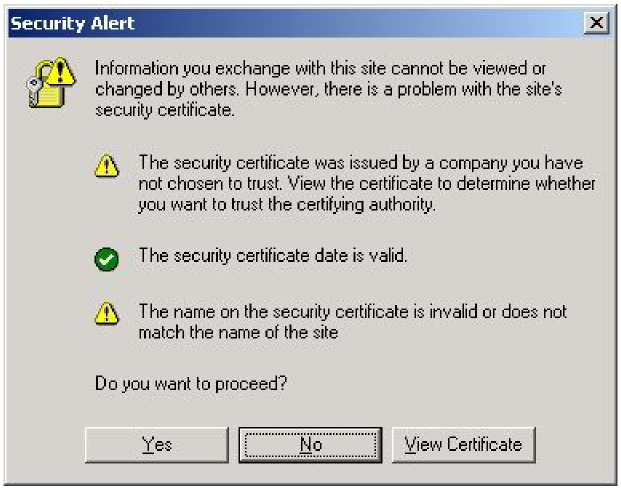

Security Alert

Information you exchange with this site cannot be viewed or changed by others. However, there is a problem with the site's security certificate.

The security certificate was issued by a company you have not chosen to trust. View the certificate to determine whether you want to trust the certifying authority.

The security certificate date is valid.

The name on the security certificate is invalid or does not match the name of the site

Do you want to proceed?

Figure 2.46: The Security Alert message reappears

12. Click **Yes** in the **Security Alert** message of Figure 2.46 to continue with the original HTTPS request for the server. The server **Welcome** page then appears. This indicates that the Certificate has been installed in Internet Explorer's certificate cache. You now need to export the certificate to the local Windows host. For that, first, follow the Tools -> Internet Options menu sequence on the web browser.

13. In the **Internet Options** dialog box that appears (see Figure 2.47), select the **Content** tab, and click on the

Certificates button within.

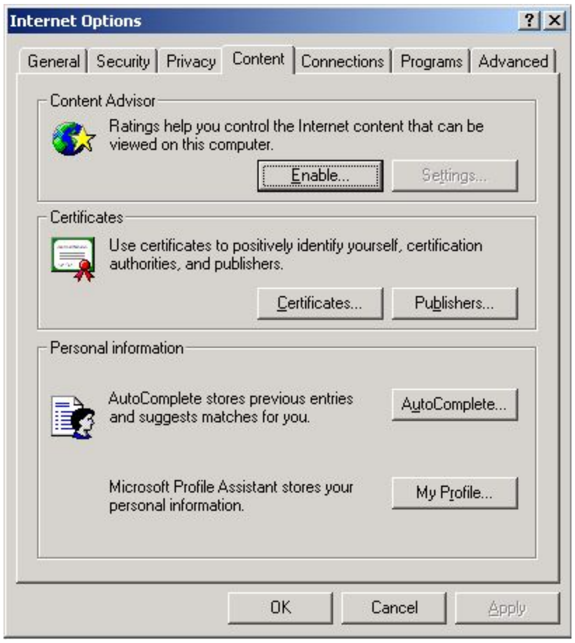

Ratings help you control the Internet content that can be viewed on this computer.

Use certificates to positively identify yourself, certification authorities, and publishers.

AutoComplete stores previous entries and suggests matches for you.

Microsoft Profile Assistant stores your personal information.

Figure 2.47: The Content tab

14. Select the **Trusted Root Certification Authorities** tab (see Figure 2.48) in the **Certificates** dialog box that appears, select the certificate that you imported from the list below, and click the **Export** button to export the certificate to the local host.

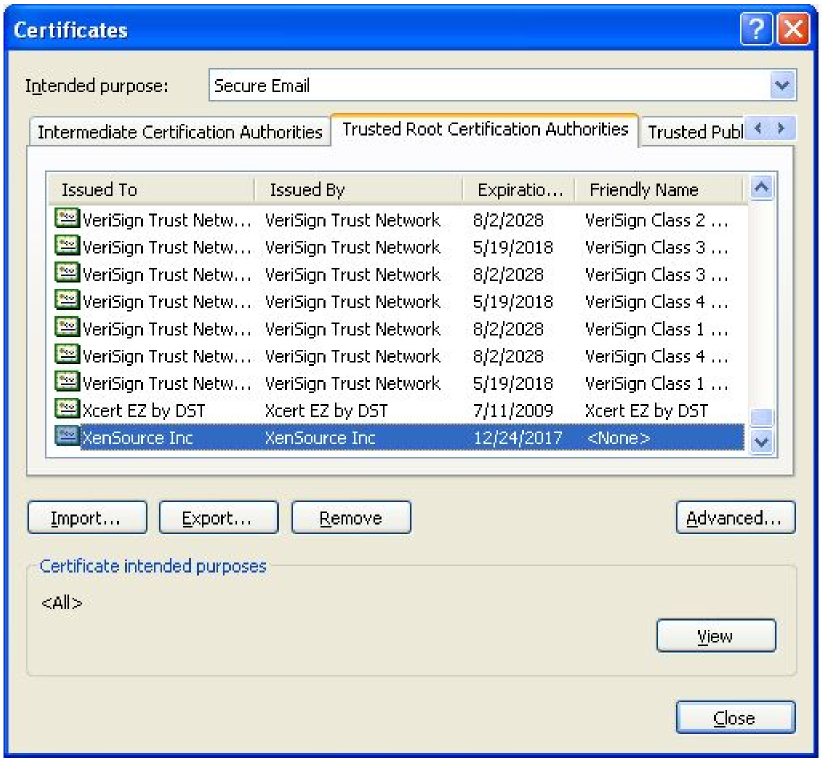

Certificates

Intended purpose: Secure Email

Intermediate Certification Authorities
Trusted Root Certification Authorities
Trusted Publ

Issued To	Issued By	Expiratio...	Friendly Name
VeriSign Trust Netw...	VeriSign Trust Network	8/2/2028	VeriSign Class 2 ...
VeriSign Trust Netw...	VeriSign Trust Network	5/19/2018	VeriSign Class 3 ...
VeriSign Trust Netw...	VeriSign Trust Network	8/2/2028	VeriSign Class 3 ...
VeriSign Trust Netw...	VeriSign Trust Network	5/19/2018	VeriSign Class 4 ...
VeriSign Trust Netw...	VeriSign Trust Network	8/2/2028	VeriSign Class 1 ...
VeriSign Trust Netw...	VeriSign Trust Network	8/2/2028	VeriSign Class 4 ...
VeriSign Trust Netw...	VeriSign Trust Network	5/19/2018	VeriSign Class 1 ...
Xcert EZ by DST	Xcert EZ by DST	7/11/2009	Xcert EZ by DST
XenSource Inc	XenSource Inc	12/24/2017	<None>

Certificate intended purposes

<All>

Figure 2.48: Viewing the Trusted Root Certification Authorities

15. The **Certificate Export Wizard** is then invoked (see Figure 2.49). Click **Next** to continue.

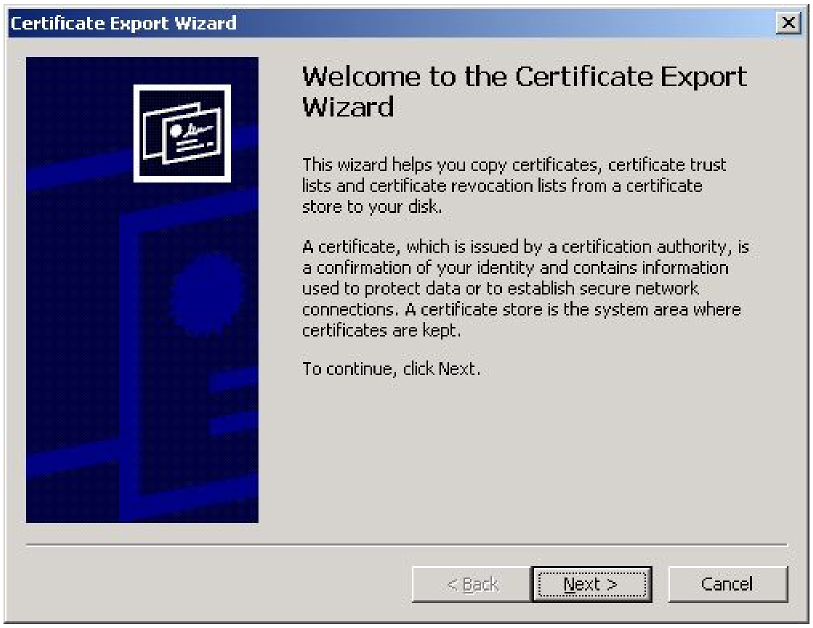

Certificate Export Wizard

Welcome to the Certificate Export Wizard

This wizard helps you copy certificates, certificate trust lists and certificate revocation lists from a certificate store to your disk.

A certificate, which is issued by a certification authority, is a confirmation of your identity and contains information used to protect data or to establish secure network connections. A certificate store is the system area where certificates are kept.

To continue, click Next.

Figure 2.49: The Certificate Export Wizard

16. When prompted to select an export file format, simply go with the default settings. Therefore, just click the **Next** button to continue (see Figure 2.50).

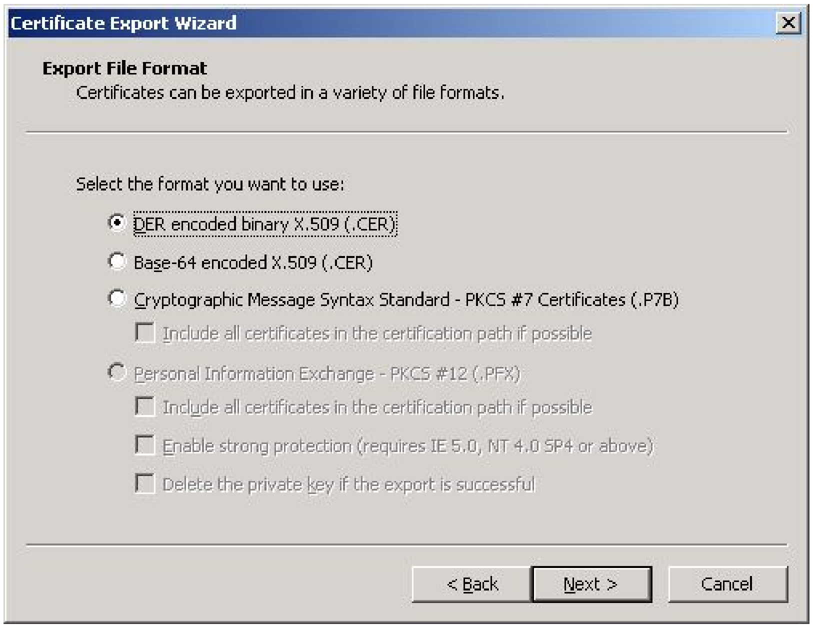

Certificate Export Wizard

Export File Format

Certificates can be exported in a variety of file formats.

Select the format you want to use:

- ● DER encoded binary X.509 (.CER)
- ○ Base-64 encoded X.509 (.CER)
- ○ Cryptographic Message Syntax Standard - PKCS #7 Certificates (.P7B)
  - □ Include all certificates in the certification path if possible
- ○ Personal Information Exchange - PKCS #12 (.PFX)
  - □ Include all certificates in the certification path if possible
  - □ Enable strong protection (requires IE 5.0, NT 4.0 SP4 or above)
  - □ Delete the private key if the export is successful

Figure 2.50: Proceeding with the default file format

17. Next, in the **File name** text box (see Figure 2.51), specify the full path to the file (on the local host) to which the certificate is to be exported. You can even click on the **Browse** button to browse for the location.

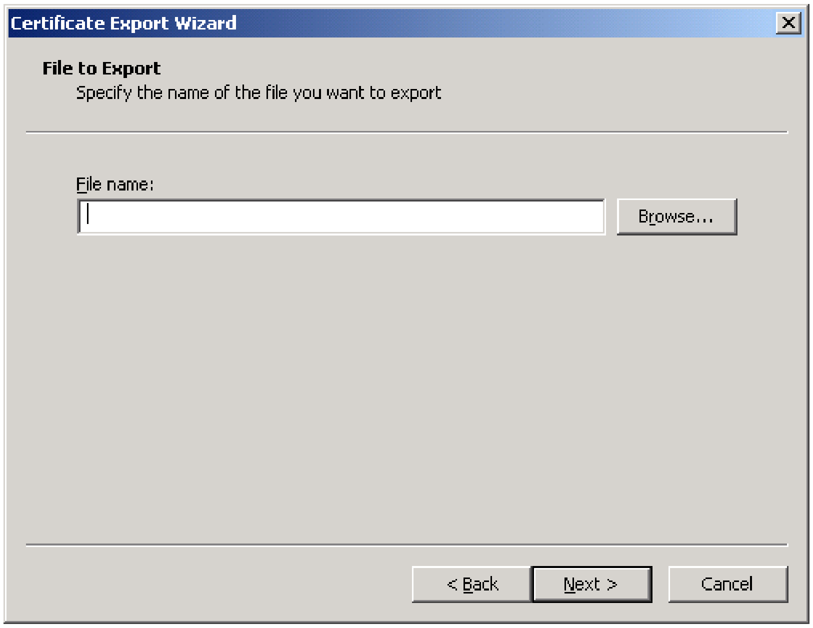

Certificate Export Wizard

File to Export

Specify the name of the file you want to export

File name:

Figure 2.51: Specifying the name of the file

18. In the **Save As** dialog box (see Figure 2.52) that appears upon clicking the **Browse** button, indicate the directory on the local host to which the certificate is to be exported, and also provide a name for the file. Then, click the **Save** button. The certificate file will be saved with the extension, **.cer**.

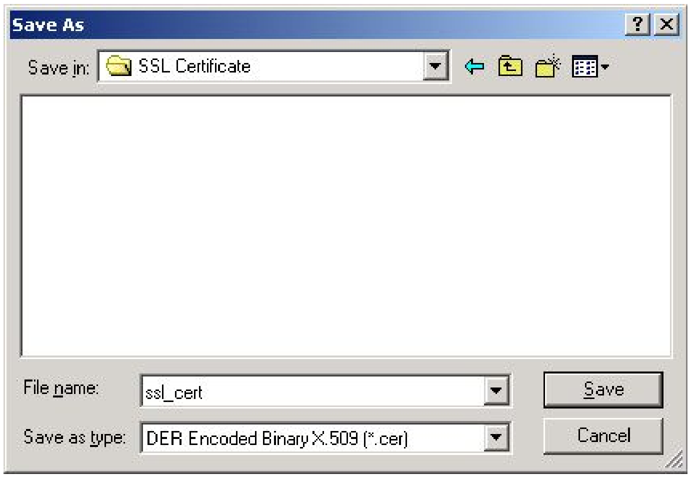

Save As

Save in: SSL Certificate

File name: ssl_cert

Save as type: DER Encoded Binary X.509 (*.cer)

Figure 2.52: Saving the certificate in a directory on the local host

19. The chosen location and specified file name will then appear against the **File name** text box (see 2.7). Click **Next** in 2.7 to continue.

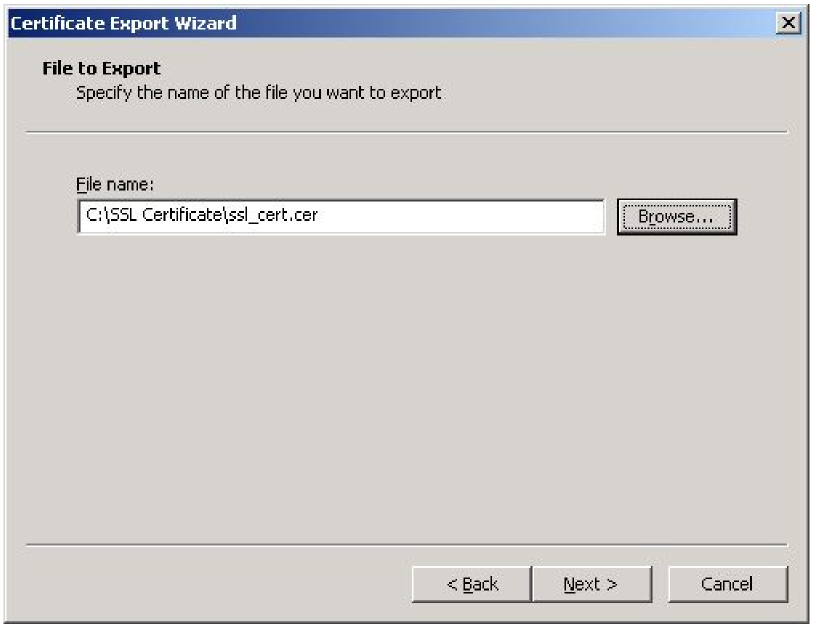

Certificate Export Wizard

File to Export

Specify the name of the file you want to export

File name:

C:\SSL Certificate\ssl_cert.cer

Figure 2.53: The chosen location appears

20. Click **Finish** in Figure 2.54 to end the export.

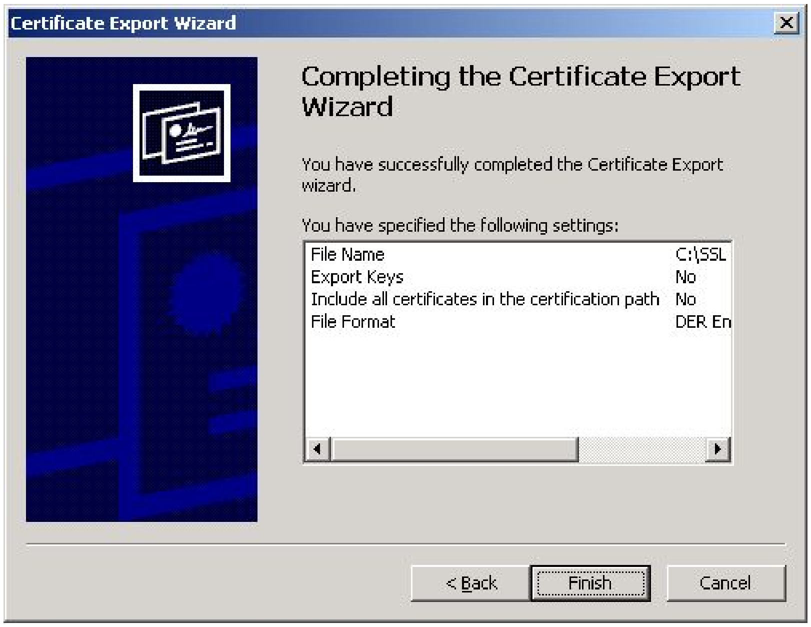

Certificate Export Wizard

Completing the Certificate Export Wizard

You have successfully completed the Certificate Export wizard.

You have specified the following settings:

File Name	C:\SSL
Export Keys	No
Include all certificates in the certification path	No
File Format	DER En

Figure 2.54: Finishing the export

2.7.1.2 Importing the Server Certificate to the Local Certificate Store of the eG Agent

The next step towards SSL-enabling the eG agent-XenServer comnmunication is to import the server certificate to the eG agent's certificate store. To do this, follow the steps given below:

1. If the eG agent is installed on the XenServer's control domain, then follow the steps given below to import the server certificate:

   - Login to the XenServer as a user with access rights to the server.
   - Open the Linux shell command.
   - Then, set the **PATH** variable appropriately i.e., **PATH = /opt/egurkha/jre/bin:$PATH**
   - Then, using the **keytool** command, import the XenServer certificate to the agent host. A sample command has been given below:

     keytool - import - file /tmp/cert/ssl_ cert.cer - alias egcert - keystore /opt/egurkha/jre/lib/security/cacerts

     The parameters expected by this command are:

     -alias : an alias name for the certificate being imported

     -file : the full path to the certificate file being imported

     -keystore : the certificate file that the JVM used by the agent checks for trusted certificates

   - This command, upon execution, will request for the keystore password. The default keystore password provided by Java is **changeit**. Provide this password and click Enter.

```
Enter keystore password: changeit
```

- Finally, the command will request your confirmation to make the imported certificate a trusted one. Enter **yes** to trust the certificate.

```
Trust this certificate? [no]:
```

- If the processing was successful, then the following message will appear.

```
Certificate was added to keystore
```

- Finally, start the agent.

2. If the eG agent is deployed on a remote Windows host, then follow the steps given below to import the certificate:

- Login to the remote agent's host.
- Open the command prompt
- Set the path to **<EG_INSTALL_DIR>\jre\bin;%path%,** using the command: **set path=<EG_ INSTALL_DIR>\jre\bin;%path%**
- Then, using the **keytool** command, import the XenServer certificate to the agent host. A sample command has been given below:

  keytool -import -file \tmp\cert\ssl_ cert.cer -alias egcert -keystore <EG_ INSTALL_ DIR>\jre\lib\security\cacerts

  The parameters expected by this command are:

  -alias :an alias name for the certificate being imported

  -file :the full path to the certificate file being imported

  -keystore : the certificate file that the JVM used by the agent checks for trusted certificates

- This command, upon execution, will request for the keystore password. The default keystore password provided by Java is **changeit**. Provide this password and click **Enter**.

```
Enter keystore password: changeit
```

- Finally, the command will request your confirmation to make the imported certificate a trusted one. Enter **yes** to trust the certificate.

```
Trust this certificate? [no]:
```

- If the processing was successful, then the following message will appear.

```
Certificate was added to keystore
```

- Finally, start the agent.

2.7.2 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:

1. Enable Password Authentication in the SSH daemon on the Linux guest; or,
2. 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:

1. Login to the Linux guest to be monitored.
2. Edit the **sshd_config** file in the **/etc/ssh** directory.
3. Check whether the **Password Authentication** flag in the **sshd_config** file is set to **no**. If so, set it to **yes**.
4. 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:

5. 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.
6. Create a directory named **.ssh** in the **<USER_HOME_DIR>** on the guest operating system, using the command: **mkdir ~/.ssh**.
7. 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.
8. Make sure that the permission of the **.ssh** directory and the **authorized_keys** file is **700**.
9. 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:

1. Login to any Linux host in your environment (even a Linux VM) as an eG user.
2. 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):
```

3. 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
```

4. 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
```

5. 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.
6. Create a directory named .ssh in the **<USER_HOME_DIR>** on the guest operating system, using the command: **mkdir ~/.ssh**.
7. 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.
8. Ensure that the **id_rsa.pub** file on the Linux guest is renamed as **authorized_keys**.
9. Repeat this procedure on every Linux guest to be monitored.
10. 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/*

11. 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:

1. 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).
2. 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.

3. 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 vSphere host. 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 vSphere/ESX host. In other words, the eG agent will be able to execute tests on all Linux guests on the target ESX host without any interruption. Therefore, press **y** to confirm key storage.

2.7.3 Troubleshooting the Failure of the eG Remote Agent to Connect to or Report Measures for Windows 2008 or Windows Vista VMs

The remote agent may not be able to connect to or collect inside view metrics from a VM running Windows 2008 or Windows Vista, if the **User Access Control** (UAC) feature is enabled on those VMs. In such a case, do the following on each of those VMs to enable the remote agent to connect to them:

1. Click **Start**
2. Type **REGEDIT**
3. Press **Enter**
4. In the left pane, browse to the following folder:

   HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\
5. Right-click on a blank area in the right pane.
6. Click **New**.
7. Click **DWORD Value**.
8. Type **LocalAccountTokenFilterPolicy**.
9. Double-click the item you just created,
10. Type **1** into the box that appears.
11. Click **OK**.
12. Restart the virtual machine.

2.7.4 Troubleshooting the Failure of the eG Remote Agent to Obtain the 'Inside View' of a Windows VM

If the eG remote agent is unable to obtain the inside view of a Windows VM, then, first check whether the agent is able to connect to the problem VM. The steps for performing this check will vary depending upon the operating system on which the remote agent executes.

Note:

The steps discussed below apply only when the following conditions are fulfilled:

a. The XenServer being monitored should be configured only with Windows VMs;

b. All the Windows VMs should belong to a single domain only - i.e., the inside view tests for the target XenServer should be configured with a single **ADMIN USER** and **ADMIN PASSWORD** only.

If the eG agent is operating on a Unix host, then the following steps will apply:

1. Login to the Unix host as eG install user.
2. Go to the shell prompt and switch to the **/opt/egurkha/lib** directory.
3. Set the **CLASSPATH** by issuing the following commands at the prompt:

   ```
   CLASSPATH=.:eg_agent.jar:eg_util.jar:$CLASSSPATH
   export CLASSPATH
   ```
4. Next, set the JRE path by issuing the following commands:

   ```
   PATH=/opt/egurkha/jre/bin:$PATH
   ```

export PATH

5. Finally, issue the following command to try connecting to the Windows VM:

   java EgWinConnect -user <username> -password <password> -domain <domainname> -IP <IP_ address>

In this command:

a. Substitute the **<username>** and **<password>** variables with the values that you have configured for the **ADMIN USER** and **ADMIN PASSWORD** parameters (respectively) for all inside view tests. While the admin user value can be easily obtained from the test configuration page in the eG administrative interface, the admin password cannot be so obtained, as it would be in an encrypted format in the user interface. To know the password therefore, do the following:

   - Open the **eg_agents.ini** file in the **/opt/egurkha/gent/config** directory on the agent host.
   - Look for entries related to the inside view tests that the agent has executed on the other Windows VMs on the target Xenserver.
   - These entries will typically include a **-admin_password** parameter, which will be followed by the decrypted admin password.
   - Make a note of the decrypted password, and substitute the **<password>** variable in the **java EgWinConnect** command with it.

b. Next, substitute the **<domainname>** variable with the value that you passed to the domain parameter in the test configuration page, and the **<IP_address>** with the IP address of the problematic Windows VM.

c. If this command fails, it is a clear indication that the remote agent is unable to communicate with the Windows VM. You will then have to investigate the reasons for the same and fix them in order to ensure that the agent is able to obtain the "inside view" of that VM.

If the eG agent is operating on a Windows host, then the following steps will apply:

1. Login to the Windows host on which the eG agent is executing.
2. Go to the command prompt and switch to the **<EG_INSTALL_DIR>\lib** directory.
3. Issue the following command to set the path to the **<EG_INSTALL_DIR>\JRE\bin** directory.

   set path=<EG_INSTALL_DIR>\JRE\bin

4. Set the **CLASSPATH** by issuing the following command at the prompt:

   set classpath=.<EG_INSTALL_DIR>\lib\eg_agent.jar;<EG_INSTALL_DIR>\lib\eg_util.jar

5. Finally, issue the following command to try connecting to the Windows VM:

   java EgWinConnect -user <username> -password <password> -domain <domainname> -IP <IP_ address>

   In this command:

   a. Substitute the **<username>** and **<password>** variables with the values that you have configured for the **ADMIN USER** and **ADMIN PASSWORD** parameters (respectively) for all inside view tests.

While the admin user value can be easily obtained from the test configuration page in the eG administrative interface, the admin password cannot be so obtained, as it would be in an encrypted format in the user interface. To know the password therefore, do the following:

- Open the **eg_agents.ini** file in the **<EG_INSTALL_DIR>\agent\config** directory on the agent host.
- Look for entries related to the inside view tests that the agent has executed on the other Windows VMs on the target Xen server.
- These entries will typically include a **-admin_password** parameter, which will be followed by the decrypted admin password.
- Make a note of the decrypted password, and substitute the **<password>** variable in the **java EgWinConnect** command with it.

b. Next, substitute the **<domainname>** variable with the value that you passed to the domain parameter in the test configuration page, and the **<IP_address>** with the IP address of the problematic Windows VM.

A sample command is given below:

C:\eGurkha\lib>java EgWinConnect - user eguser - password C13120CB9E5D4B1423419897808BAE65 -domain mas -ip 192.168.10.216

6. If the command is successful, the following output will appear:

```
********************************************************************************
Attempt to connect and execute for 192.168.10.216
Output is [[Ok, , , , , Windows IP Configuration, , , , , , , , , Ethernet adapt
er Local Area Connection 2:, , , , , ,    Connection-specific DNS Suffix  . : ,
, ,    IP Address. . . . . . . . . . . . : 192.168.10.216, , ,    Subnet Mask .
. . . . . . . . . . : 255.255.255.0, , ,    Default Gateway . . . . . . . . . :
192.168.10.2, , , , EgDone 0(0x0)], []
********************************************************************************
```

7. The command may fail when it encounters one of the following errors. The reasons for these errors and the recommended resolution for the same have been provided below.

Error	Reason	Fix
Counldn't connect to \\<IP_of_Windows_VM>\ADMIN$ **Logon failure: unknown user name or bad password**	Occurs if the inside-view tests have been configured with an incorrect **DOMAIN**, invalid **ADMIN USER** name, or a wrong **ADMIN PASSWORD**.	Reconfigure the tests with the valid credentials of a domain administrator
Counldn't connect to \\<IP_of_Windows_VM>\ADMIN$	• Can occur if the **ADMIN$** share has not been enabled	• Enable the ADMIN$ share on the target VM.

Error	Reason	Fix
The network path was not found.	on the target Windows VM; - Can occur if the Windows firewall is blocking connection to the VM, or if File/Print Sharing has not been enabled yet.	- Provide domain administrator with full access to the ADMIN$ share. - Try connecting to the Windows VM remotely; if the problem persists: - Reconfigure the Windows firewall to allow communication between the remote agent and the Windows VM; - Reconfigure the Windows firewall to allow File/Print Sharing
Couldn't copy service to \\<IP_of_Windows_VMs>\ADMIN$ **Access is denied.**	- Can occur if the **ADMIN$** share exists but the domain administrator does not have permission to access the shared folder; - Can occur if the **ADMIN$** share exists, but the user with full access to the shared folder is not the domain administrator.	Provide domain administrator with full access to the **ADMIN$** share.

2.7.5 Troubleshooting the Failure of eG Agent to Connect to the Xen Server and Collect Metrics from it

Sometimes, if the eG agent is unable to establish a connection with the Xen Server, no tests can be run on the server. Similarly, if the eG tests are incorrectly configured, the tests will not execute in the desired manner, nor will the required metrics be reported.

To enable administrators to quickly isolate issues in connection and test configuration, the eG agent is bundled with a **EgXenReport** command. Upon execution, this command reports real-time statistics that shed light on connection bottlenecks between the agent and the XenServer, improper test configuration, inexplicable delays in connection, and more. The procedure for executing this command may slightly differ depending upon the operating system on which the eG agent is executing. The OS-specific procedures have been detailed below.

Executing the 'EgXenReport' command on a Windows host

1. Login to the Windows system hosting the eG agent.
2. Go to the command prompt.

3. Set the Java **path** variable using the following command:

   set path=.;<EG_INSTALL_DIR>\JRE\bin;%path%

   Example: **set path=.;C:\eGurkha\JRE\bin;%path%**

4. Set the **classpath** variable using the following command:

   set classpath=.;< EG_ INSTALL_ DIR >\lib\ws- commons- util.jar;<eGurkha_ Install_ dir>\lib\XenSourceAPI.jar< EG_ INSTALL_ DIR >\lib\xmlrpc- client.jar; <EG_ INSTALL_ DIR>\lib\xmlrpc-common.jar; <EG_INSTALL_DIR>\lib\eg_agent.jar;<eGurkha_Install_dir>\lib\eg_ util.jar;%classpath%

   Example: **set classpath=.;C:\eGurkha\lib\ws- commons- util.jar;C:\eGurkha\lib\XenSourceAPI.jar;C:\eGurkha\lib\xmlrpc- client.jar;C:\eGurkha\lib\xmlrpc-common.jar;C:\eGurkha\lib\eg_agent.jar;C:\eGurkha\lib\eg_util.jar;%classpath%**

5. Run the **EgXenReport** command as indicated below:

   java EgXenReport <ipaddress> <port> <user> <password> <ssl>

   Here:
   <ipaddress> refers to the IP address of the managed XenServer in the eG Enterprise system

   <port> refers to the port number that was defined when adding the XenServer to the eG Enterprise system

   <user> refers to the user name configured against the **XEN USER** test parameter

   <password> refers to the password that you have specified against the **XEN PASSWORD** test parameter. This password should be provided in a decrypted format. For the decrypted password, do the following:

   - Open the **eg_agents.ini** file in the **<EG_INSTALL_DIR>\agent\config** directory on the agent host.
   - Look for entries related to the tests that the agent has executed on the target Xen server.
   - These entries will typically include a **-xen_password** parameter, which will be followed by the decrypted xen password.
   - Make a note of the decrypted password, and substitute the **<password>** variable in the **java EgXenReport** command with it.

   <ssl> indicates the status of the **SSL** flag in the test configuration (i.e., whether **true/false**)

   Example: **java EgXenReport 192.168.10.156 80 root C13120CB9E5D4B1423419897808BAE65 false**

6. Figure 2.55 displays the sample output of the command.

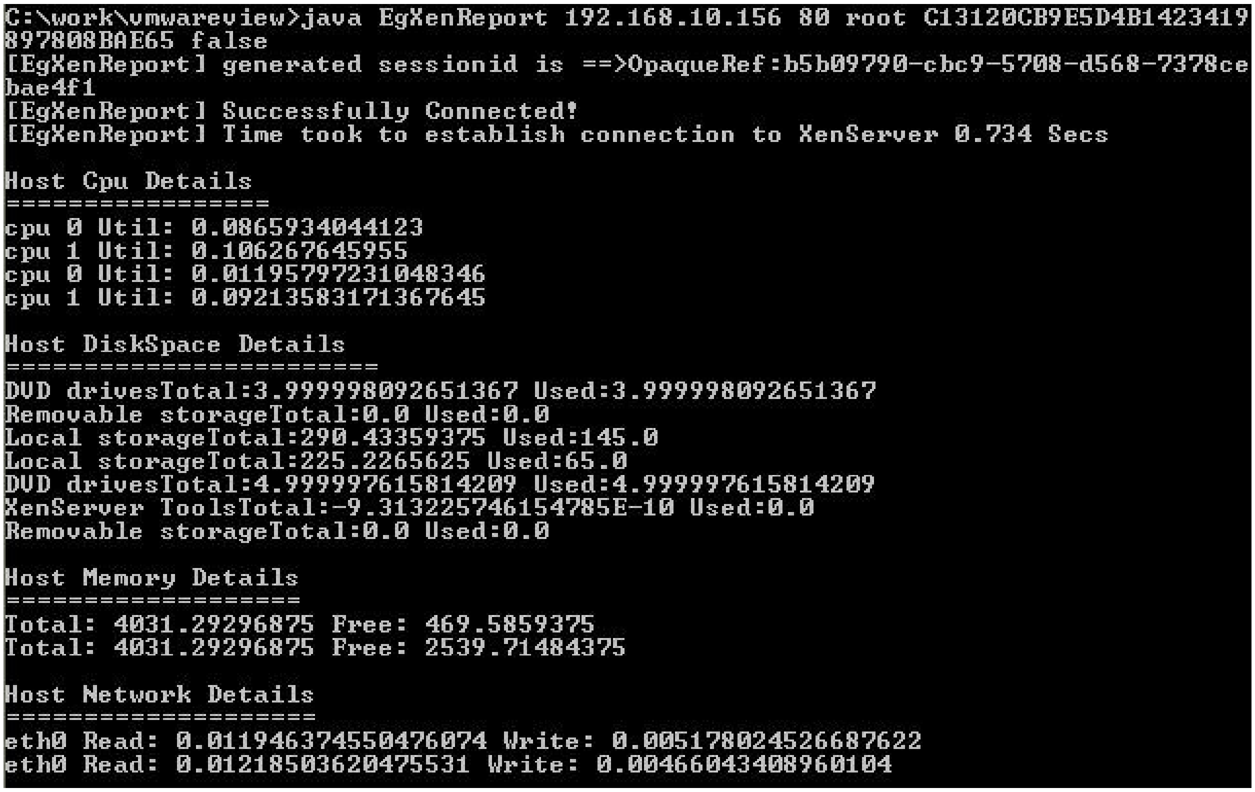

```
C:\work\vmwareview>java EgXenReport 192.168.10.156 80 root C13120CB9E5D4B1423419
897808BAE65 false
[EgXenReport] generated sessionid is ==>OpaqueRef:b5b09790-cbc9-5708-d568-7378ce
bae4f1
[EgXenReport] Successfully Connected!
[EgXenReport] Time took to establish connection to XenServer 0.734 Secs

Host Cpu Details
=================
cpu 0 Util: 0.0865934044123
cpu 1 Util: 0.106267645955
cpu 0 Util: 0.0119579723104834б
cpu 1 Util: 0.09213583171367645

Host DiskSpace Details
========================
DVD drivesTotal:3.999998092651367 Used:3.999998092651367
Removable storageTotal:0.0 Used:0.0
Local storageTotal:290.43359375 Used:145.0
Local storageTotal:225.2265625 Used:65.0
DVD drivesTotal:4.999997615814209 Used:4.999997615814209
XenServer ToolsTotal:-9.313225746154785E-10 Used:0.0
Removable storageTotal:0.0 Used:0.0

Host Memory Details
===================
Total: 4031.29296875 Free: 469.5859375
Total: 4031.29296875 Free: 2539.71484375

Host Network Details
====================
eth0 Read: 0.011946374550476074 Write: 0.005178024526687622
eth0 Read: 0.01218503620475531 Write: 0.00466043408960104
```

Figure 2.55: Sample output of the EgXenReport command if connection succeeds

7. If the command is able to connect to the XenServer successfully, the output will indicate the same with the message **Successfully Connected**. In addition, the time taken for this connection to be established will also be reported as part of the output. With the help of this metric, you can quickly detect connection latencies/bottlenecks. Besides the above, the command collects a wide variety of host-level metrics from the XenServer in real-time and reports the same in the output.

8. On the other hand, if the command is unable to establish a connection with the XenServer, the connection failure and the reasons for the same will be cited as part of the output (see Figure 2.56).

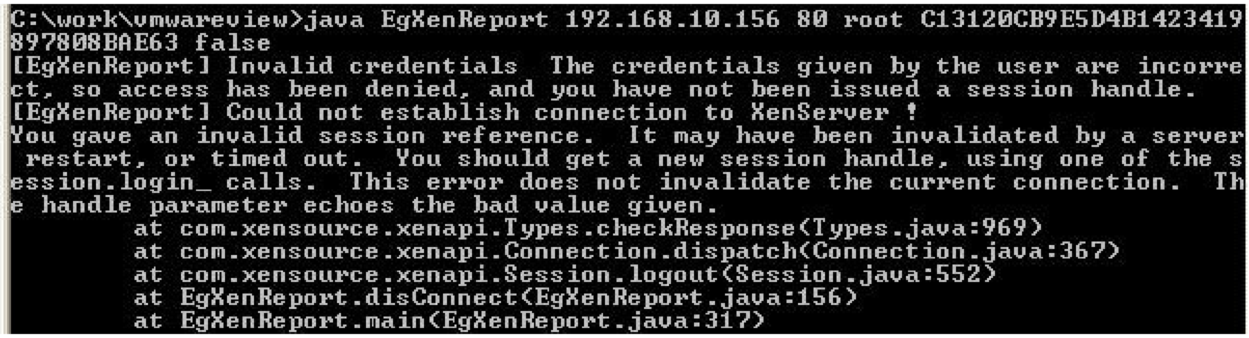

```
C:\work\vmwareview>java EgXenReport 192.168.10.156 80 root C13120CB9E5D4B1423419
897808BAE63 false
[EgXenReport] Invalid credentials  The credentials given by the user are incorre
ct, so access has been denied, and you have not been issued a session handle.
[EgXenReport] Could not establish connection to XenServer !
You gave an invalid session reference.  It may have been invalidated by a server
 restart, or timed out.  You should get a new session handle, using one of the s
ession.login_ calls.  This error does not invalidate the current connection.  Th
e handle parameter echoes the bad value given.
        at com.xensource.xenapi.Types.checkResponse(Types.java:969)
        at com.xensource.xenapi.Connection.dispatch(Connection.java:367)
        at com.xensource.xenapi.Session.logout(Session.java:552)
        at EgXenReport.disConnect(EgXenReport.java:156)
        at EgXenReport.main(EgXenReport.java:317)
```

Figure 2.56: Sample output of the EgXenReport command if connection fails

If the connection failed because incorrect values were provided as arguments to the **EgXenReport** command, then the output will indicate the exact argument that caused the failure. This way, you will be able to check the correctness of your test configuration, spot errors in the same, and change it.

Executing the 'EgXenReport' command on a Solaris/Linux host

1. Login to the Solaris/Linux system hosting the eG agent.
2. Go to the command prompt.
3. Set the Java **path** variable using the following command:

   PATH=.:<EG_INSTALL_DIR>/JRE/bin:$PATH

   export $PATH

   Example:

   PATH=.:/opt/egurkha/jre/bin:$PATH

   export $PATH

4. Set the **classpath** variable using the following command:

   CLASSPATH=.:< EG_ INSTALL_ DIR >/lib/ws- commons- util.jar:< EG_ INSTALL_ DIR >/lib/XenSourceAPI.jar:< EG_ INSTALL_ DIR >/lib/xmlrpc- client.jar:< EG_ INSTALL_ DIR>/lib/xmlrpc-common.jar:<EG_INSTALL_DIR>/lib/eg_agent.jar:<EG_INSTALL_DIR>/lib/eg_ util.jar:$CLASSPATH

   export $CLASSPATH

   Example:

   CLASSPATH=.:/opt/egurkha/lib/ws- commons- util.jar:/opt/egurkha/lib/XenSourceAPI.jar:/opt/egurkha/lib/xmlrpc- client.jar:/opt/egurkha/lib/xmlrpc- common.jar:/opt/egurkha/lib/eg_ agent.jar:/opt/egurkha/lib/eg_ util.jar:$CLASSPATH

   export $CLASSPATH

5. Run the **EgXenReport** command as indicated below:

   java EgXenReport <ipaddress> <port> <user> <password> <ssl>

   Here:

   <ipaddress> refers to the IP address of the managed XenServer in the eG Enterprise system

   <port> refers to the port number that was defined when adding the XenServer to the eG Enterprise system

   <user> refers to the user name configured against the **XEN USER** test parameter

   <password> refers to the password that you have specified against the **XEN PASSWORD** test parameter. This password should be provided in a decrypted format. For the decrypted password, do the following:

   - Open the **eg_agents.ini** file in the **/opt/egurkha/agent\config** directory on the agent host.
   - Look for entries related to the tests that the agent has executed on the target Xen server.
   - These entries will typically include a **-xen_password** parameter, which will be followed by the decrypted xen password.

- Make a note of the decrypted password, and substitute the **<password>** variable in the **java EgXenReport** command with it.

<ssl> indicates the status of the **SSL** flag in the test configuration (i.e., whether **true/false**)

Example: **java EgXenReport 192.168.10.156 80 root C13120CB9E5D4B1423419897808BAE65 false**

6. The sample output of the command and their interpretations are available in the Figure 2.56.

The Citrix XenServer - VDI Monitoring Model

Using the *Citrix XenServer - VDI* monitoring model (see Figure 3.1) that eG Enterprise offers, you can monitor the overall health of the Citrix XenServer and its virtual desktops, the number of users interacting with each desktop, and the quality of every user's experience with the desktop.

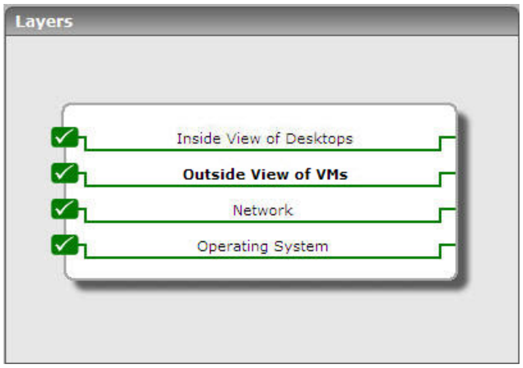

Figure 3.1: Figure 3.1: The layer model of Citrix XenServer - VDI

Each layer of Figure 3.1 is mapped to tests that report useful statistics, which enable administrators to figure out the following:

- How many desktops are powered on simultaneously on the XenServer?
- How much CPU, memory, disk and network resources is each desktop taking?
- What is the typical duration of a user session?
- Which users are logged on and when did each user login?
- Who has the peak usage times?
- What applications are running on each desktop?

Since the **Operating System** and **Network** layers of Figure 3.1 have already been discussed in Chapter 2 of this document, the sections that follow will elaborate on the **Xen Guests** and **Virtual Desktop** layer only.

3.1 The Outside View of VMs Layer

The **Outside View of VMs** layer provides the host operating system's view of the resource usage levels of each of the desktops hosted on it. Using the information reported by this test, administrators can:

- Determine which of the desktops are 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 desktops is receiving a very high rate of requests compared to the others, this desktop may be a candidate for migration to another XenServer, so as to minimize the impact it has on the other desktops on the current server.

- Determine times when sudden or steady spikes in the physical resource utilization are caused by the desktops
- Know which guest systems at what times experienced heavy session loads or unexpected session logouts

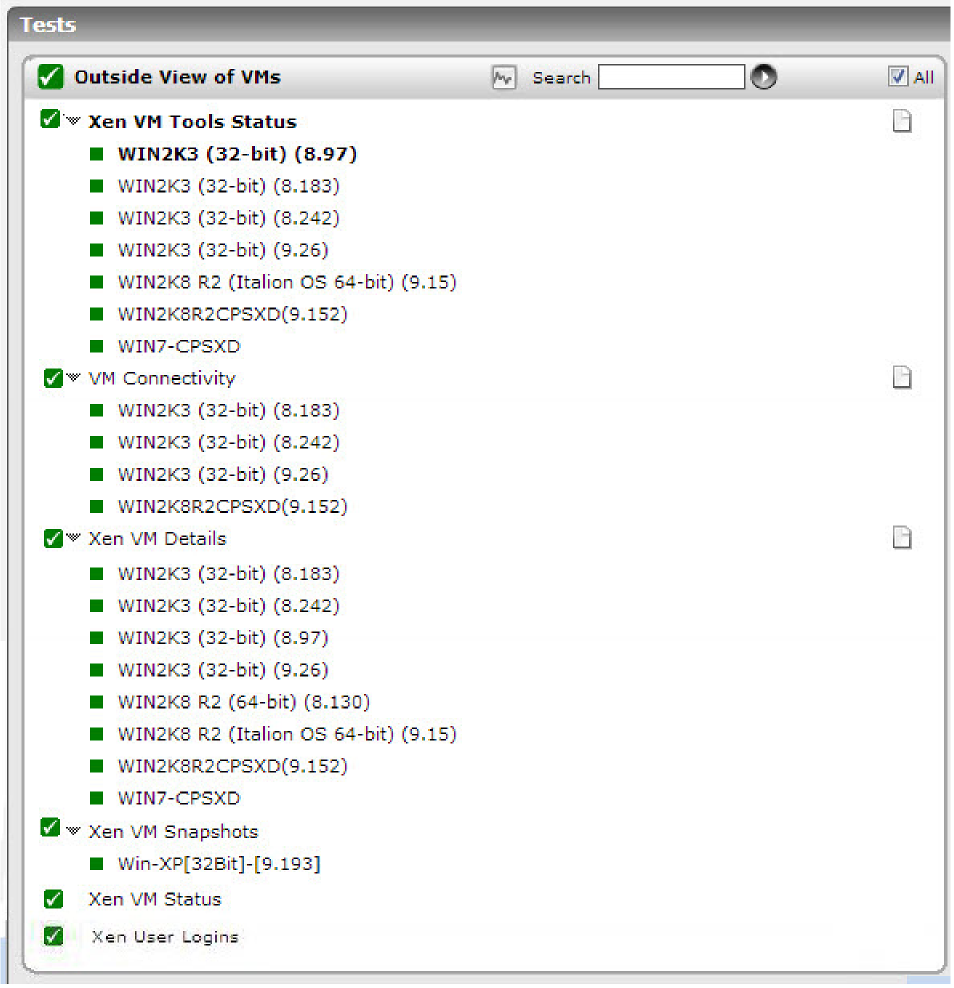

Tests

☑ **Outside View of VMs** Search ☑ All

☑ **Xen VM Tools Status**
- 🟩 **WIN2K3 (32-bit) (8.97)**
- 🟩 WIN2K3 (32-bit) (8.183)
- 🟩 WIN2K3 (32-bit) (8.242)
- 🟩 WIN2K3 (32-bit) (9.26)
- 🟩 WIN2K8 R2 (Italion OS 64-bit) (9.15)
- 🟩 WIN2K8R2CPSXD(9.152)
- 🟩 WIN7-CPSXD

☑ VM Connectivity
- 🟩 WIN2K3 (32-bit) (8.183)
- 🟩 WIN2K3 (32-bit) (8.242)
- 🟩 WIN2K3 (32-bit) (9.26)
- 🟩 WIN2K8R2CPSXD(9.152)

☑ Xen VM Details
- 🟩 WIN2K3 (32-bit) (8.183)
- 🟩 WIN2K3 (32-bit) (8.242)
- 🟩 WIN2K3 (32-bit) (8.97)
- 🟩 WIN2K3 (32-bit) (9.26)
- 🟩 WIN2K8 R2 (64-bit) (8.130)
- 🟩 WIN2K8 R2 (Italion OS 64-bit) (9.15)
- 🟩 WIN2K8R2CPSXD(9.152)
- 🟩 WIN7-CPSXD

☑ Xen VM Snapshots
- 🟩 Win-XP[32Bit]-[9.193]

☑ Xen VM Status

☑ Xen User Logins

Figure 3.2: Figure 3.2: The tests mapped to the Xen VMs layer

The sections that follow will discuss the Xen VM Status test, Xen VM Details test, and the Xen User Logins test only.

3.1.1 VM Details - Xen Test

This test monitors the amount of the physical server's resources that each guest on an XenServer 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 taking up the maximum memory utilization, which guest has the maximum disk activity, 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. Also, virtual guests can be organized into resource pools, and allocation of resources can be made at the resource-pool level. In this case, virtual guests allocated to the same resource pool contend for the resources allocated to the resource pool.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each guest configured on the XenServer host

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer’s configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG

agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

- Obtain the server-certificate for the XenServer
- Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **AGGREGATE USER SESSIONS** -This flag is closely related to the **REPORT BY USER** flag. Since the **REPORT BY USER** flag is set to **No** by default for a Citrix XenServer, this test will, by default, ignore the status of the **AGGREGATE USER SESSIONS** flag while monitoring that server. In case of the Citrix XenServer - VDI on the other hand, the **REPORT BY USER** flag is set to **Yes** by default. Therefore, the status of the **AGGREGATE USER SESSIONS** flag gains significance in the case of this model. By default, the **AGGREGATE USER SESSIONS** flag is set to **No**. This implies that if a single user is currently logged into multiple guests, then this test, by default, will report a set of measures for every *username* on *guestname*. On the other hand, if the status of this flag is changed to **Yes**, then, this test will report a set of (aggregated) measures for every distinct user to the virtual desktop environment. In other words, this test will report measures that are aggregated across all the currently active sessions for a user, spanning multiple VMs.

15. **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 *virtualmachine* 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.

16. **REPORT POWERED ON** - You can set the **REPORT POWERED ON** status to **Yes**, so that the test reports an additional measure, Is VM powered on?, revealing whether a guest OS is currently running or not. The default status of this flag is set to Yes for a Citrix XenServer. For a Citrix XenServer – VDI component on the other hand, this flag is set to No by default. This is because, in such environments, the virtual desktops will be in the powered-off state most of the time.

17. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports *80* or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

18. **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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current sessions:	This measure is relevant only for monitoring of virtual desktops (i.e., for Citrix XenServer - VDI servers). When reporting metrics for specific users, this metric indicates the number of sessions that each user has currently logged into; this measure will be available only if the test reports measures per currently logged in user.	Number	This is a good indicator of how busy the user is. The detailed diagnosis of this measure, if enabled, reveals the guests to which the user is currently logged on to.
Is VM powered on?:	Whether the virtual machine is currently running on the XenServer host or no.		While the test reports a wide variety of other metrics too for virtual machines that are alive, only the powered on status is indicated for virtual machines that are currently not available. The value Yes for this measure indicates that the guest is up and running. The value No could indicate that the guest has been powered-off; it could also indicate that XenMotion has moved the guest to a different server. The numeric values that correspond to each of the powered- on states

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>discussed above are listed in the table below:<table><tr><th>State</th><th>Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table><b>Note:</b>
By default, this measure reports <i>Yes</i> or <i>No</i> to indicate VM status. The graph of this measure however, represents the status of a VM using the numeric equivalents - 0 or 1.</td></tr>
<tr><td><b>Physical CPU usage:</b></td><td>Indicates the percentage of physical CPU used by the guest.</td><td>Percent</td><td>A high value for this measure indicates a virtual machine that is using a lot of the processor - possibly because one or more processes on this VM are taking a lot of CPU.</td></tr>
<tr><td><b>Free physical memory:</b></td><td>Indicates the amount of memory available for use with the guest.</td><td>MB</td><td>Ideally, this value should be high. A low or consistent decrease in this value denotes that the application (s) executing on the guest are consuming memory excessively. You might want to consider increasing the memory allocated to the guest. XenServer Enterprise and XenServer Standard allow that a Linux/Windows VM can use up to 32GB of memory. Moreover, Xen has implemented a balloon driver concept for each domain, enabled independently, that allows the operating system to adjust its current memory allocation up to the maximum limit configured. This allows “unused” allocation to be consumed in other areas, potentially allowing for stable over-commitment of memory resources. Because of this constantly changing</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			memory allocation, memory is allocated and freed dynamically at a granularity of the page-level.
Total physical memory allocated:	Indicates the amount of physical memory currently allocated to the guest.	MB	
Used physical memory:	Indicates the amount of memory used by the guest.	MB	
Usage of allocated memory:	Indicates the percentage of allocated memory that is being used by the guest.	Percent	High memory consumption over long periods can deplete the free memory on the guest, causing prolonged delays in the execution of the application (s) hosted by the guests.
Disk capacity:	Indicates the total allocated disk space of the guest.	MB	
Disk read rate:	Indicates the rate at which the guest read from the disk.	Kbytes/Sec	
Disk write rate:	Indicates the rate at which the guest wrote data to the disk.	Kbytes/Sec	
Network data received:	Indicates the network I/O reads performed by the guest.	Mbps	
Network data transmitted:	Indicates the network I/O writes performed by the guest.	Mbps	
Virtual CPU usage:	Indicates the percentage of allocated CPU resources that this VM is currently	Percent	Comparing the value of this measure across VMs will enable you to accurately identify the VMs on which

Measurement	Description	Measurement Unit	Interpretation
	using.		CPU- intensive applications are executing.

3.1.2 VM Status – Xen Test

This test enables administrators to determine how many guests have registered with the XenServer, and how many of these are currently running. In addition, the test also indicates whether any guests have migrated to or from the virtual server. XenMotion is a feature using which virtual machines can be moved from server to server within a Citrix XenServer resource pool without service interruption, making zero-downtime server maintenance possible, and enabling administrators to move resource-hungry running applications to take advantage of available compute power. In short, the Xen VM Status test monitors the effectiveness of XenMotion.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for the XenServer being monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.

- Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.
8. **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*.
9. **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 **2.7.2**.
- **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 PASSWORD**s 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 **2.5.1.1**.
- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.
12. **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.
13. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.
14. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to *1:1*, it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Added guests:	Indicates the number of guests that were newly added to the XenServer during this measurement period.	Number	The detailed diagnosis of these measures, if enabled, lists the virtual machines that were migrated by XenMotion to or from (as the case may be) the XenServer.
Removed guests:	Indicates the number of guests that were newly removed from the XenServer during this measurement period.	Number	
Registered guests:	Indicates the total number of virtual machines that have been registered with the XenServer.	Number	
Running guests:	Indicates the total number of virtual machines that are currently running.	Number	
Halted guests:	Indicates the number of guests that were stopped.	Number	
Suspended guests:	Indicates the total number of virtual machines that have been suspended.	Number	You can suspend a XenVM, saving its state to a file. This allows you to shut down the physical XenServer host and later, after rebooting it, return the XenVM to its original running state.
Guests with users:	Indicates the number of powered on guests with users logged in.	Number	To know which guests the users have logged into, use the detailed diagnosis capability of this measure (if enabled).
Guests without users:	Indicates the number of powered on guests without any users logged in.	Number	

The detailed diagnosis of the *Registered guests* measures, if enabled, lists the virtual machines that are currently registered with the XenServer (see Figure 3.3).

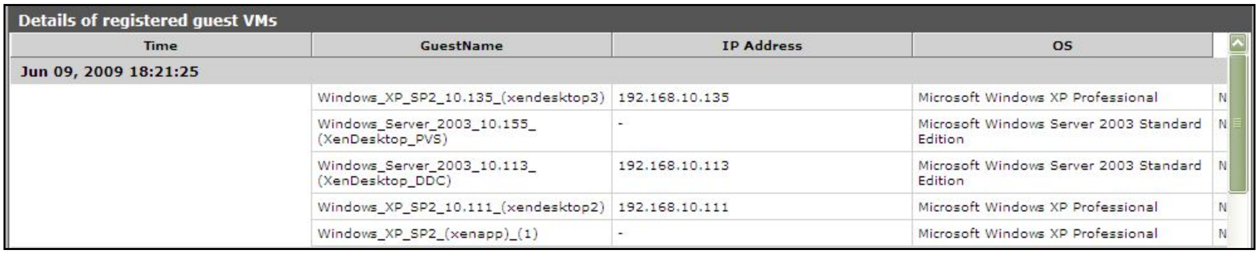

Details of registered guest VMs

Time	GuestName	IP Address	OS	
Jun 09, 2009 18:21:25				
	Windows_XP_SP2_10.135_(xendesktop3)	192.168.10.135	Microsoft Windows XP Professional	N
	Windows_Server_2003_10.155_ (XenDesktop_PVS)	-	Microsoft Windows Server 2003 Standard Edition	N
	Windows_Server_2003_10.113_ (XenDesktop_DDC)	192.168.10.113	Microsoft Windows Server 2003 Standard Edition	N
	Windows_XP_SP2_10.111_(xendesktop2)	192.168.10.111	Microsoft Windows XP Professional	N
	Windows_XP_SP2_(xenapp)_(1)	-	Microsoft Windows XP Professional	N

Figure 3.3: The detailed diagnosis of the Registered guests measure

The detailed diagnosis of the *Halted guests* measure, if enabled, provides the IP addresses of the guests that stopped running, the guest names, and the operating system on which the guest is executing (see Figure 3.4).

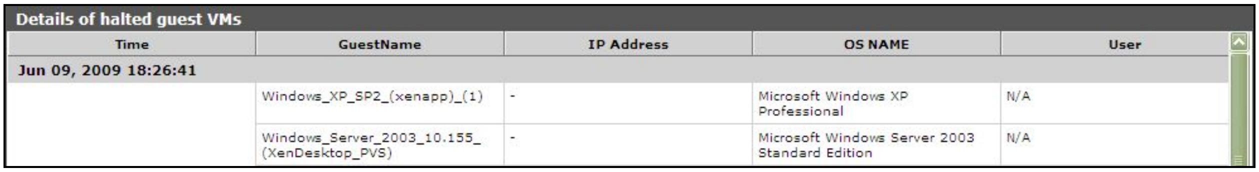

Details of halted guest VMs

Time	GuestName	IP Address	OS NAME	User
Jun 09, 2009 18:26:41				
	Windows_XP_SP2_(xenapp)_(1)	-	Microsoft Windows XP Professional	N/A
	Windows_Server_2003_10.155_ (XenDesktop_PVS)	-	Microsoft Windows Server 2003 Standard Edition	N/A

Figure 3.4: The detailed diagnosis of the Halted guests measure

The detailed diagnosis of the *Running guests* measure, if enabled, provides details pertaining to VMs that are currently running.

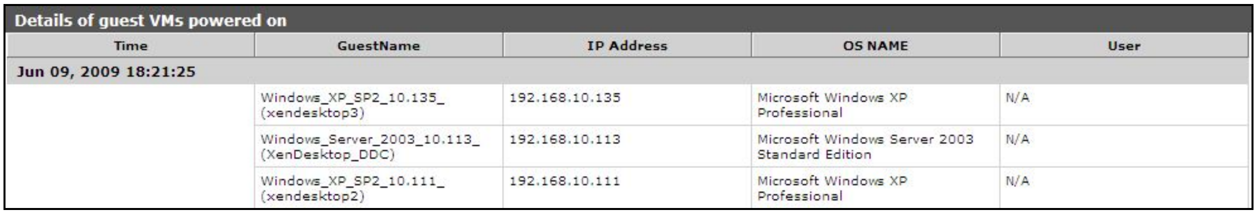

Details of guest VMs powered on

Time	GuestName	IP Address	OS NAME	User
Jun 09, 2009 18:21:25				
	Windows_XP_SP2_10.135_ (xendesktop3)	192.168.10.135	Microsoft Windows XP Professional	N/A
	Windows_Server_2003_10.113_ (XenDesktop_DDC)	192.168.10.113	Microsoft Windows Server 2003 Standard Edition	N/A
	Windows_XP_SP2_10.111_ (xendesktop2)	192.168.10.111	Microsoft Windows XP Professional	N/A

Figure 3.5: The detailed diagnosis of the Running guests measure

3.1.3 User Logins – Xen Test

This test monitors the user logins to guests and reports the total count of logins and logouts.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for XenServer being monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.

3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.

4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.

6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.

7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

   - **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

14. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to *1:1*, it means that detailed measures will be generated every

time this test runs, and also every time the test detects a problem.

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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current sessions:	Indicates the number of user sessions that are currently active across all guests	Number	This is a good indicator of the session load on the guests.
New logins:	Indicates the number of new logins to the guests.	Number	A consistent zero value could indicate a connection issue. If this measure reports a non-zero value, use the detailed diagnosis of the measure to know which user logged into which VM, when.
Percent new logins:	Indicates the percentage of current sessions that logged in during the last measurement period.	Percent	
Sessions logging out:	Indicates the number of sessions that logged out.	Number	If all the current sessions suddenly log out, it indicates a problem condition that requires investigation. The detailed diagnosis of this measure lists the sessions that logged out.

3.1.4 Xen VM Snapshots Test

XenServer provides a convenient snapshotting mechanism that can take a snapshot of a VM storage and metadata at a given time. Where necessary, IO is temporarily halted while the snapshot is being taken to ensure that a selfconsistent disk image can be captured. Snapshot operations result in a snapshot VM that is similar to a template. The VM snapshot contains all the storage information and VM configuration, including attached VIFs, allowing them to be exported and restored for backup purposes.

Though snapshot files are small in size initially, they will grow as writes are made to the VM's disk files. If the number and size of the snapshot files grow significantly over time, they might end up eroding considerable storage space, thereby choking VM operations. To conserve space, administrators need to continuously track snapshot growth per VM, identify 'heavy-weight' snapshots that may not be of use any longer, and purge them. The **Xen VM Snapshots** test helps administrators achieve the same. While the measures reported by the test capture the snapshot file count per VM and the total size of the snapshot files of a VM, the detailed diagnosis reveals the size of each snapshot, thus enabling administrators to quickly spot those snapshot files that are too large in size.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each VM on the target XenServer host

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right

to connect to the XenServer console via SSH.

- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.
9. **AGELIMIT** - By default, 15 days is set as **AGELIMIT**. This implies that the test will report all those snapshots that are more than 15 days old as **Old snapshots**. If required, you can change the **AGELIMIT**.
10. **SIZELIMIT**- By default, 10000 KB is set as the **SIZELIMIT**. This implies that the test will report all those snapshots that have a size more than 10000 KB as **Large snapshots**. If required, you can change the **SIZELIMIT**.
11. **SHOW TOP** – By default, the detailed diagnosis of the *Number of snapshots* measure provides details of all the snapshots on a VM. Likewise, the detailed diagnosis of the *Large snapshot count* and *Aged snapshot count* measures provides the details of all large-sized snapshots and old snapshots, respectively. This is why, the **SHOW TOP** parameter is set to *All* by default. To make sure that the detailed diagnosis of the each of these measures lists, for instance, only the top-5 snapshots in terms of size/age (as the case may be), then, specify **5** against **SHOW TOP**. If this is done, then:

    a. The detailed diagnosis of the *Number of snapshots* and *Aged snapshot count* measures will provide the details of only the top-5 snapshots in terms of age.

b. The detailed diagnosis of the *Large snapshot count* measure will provide the details of only the top-5 snapshots in terms of size.

In the same way, you can specify any non-zero value against show top to view only that many top

12. **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.

13. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to *1:1*, it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.

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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Number of snapshots:	Indicates the number of snapshot files of this VM that are currently available.	Number	A number of snapshots of a VM provides administrators with multiple restore points. On the flipside though, a high number of snapshots can also be considered a waste of valuable disk space, especially if many of the snapshots hold less critical, but heavy-

Measurement	Description	Measurement Unit	Interpretation
			weight changes/writes to the disk. To accurately identify those snapshots that are consuming disk space excessively, and to learn when they were created, who their parents are, and their current consistent file system state, use the detailed diagnosis of this measure.
Aged snapshots count:	Indicates the number of snapshots that are of an age over the configured **AGELIMIT**.	Number	Use the detailed diagnosis of this measure to identify the old snapshots, so that you can figure out whether they deserve to be retained or not. While many snapshots provide essential restore points for VMs, many others hold less critical information. The 'less useful' snapshots can be eliminated to save disk space.
Large snapshots count:	Indicates the number of snapshots that are of a size more than the configured **SIZELIMIT**.	Number	Though small in size initially, snapshots can grow with time, but can never grow beyond the original disk file size. If a marked increase is noticed in the value of this measure over time, it could indicate that a number of snapshots are rapidly growing in size. To know which snapshots are growing beyond the size limit set, use the detailed diagnosis of this measure.

To accurately identify those snapshots that are consuming disk space excessively, and to learn when they were created, who their parents are, and their current consistent file system state, use the detailed diagnosis of the *Number of snapshots* measure.

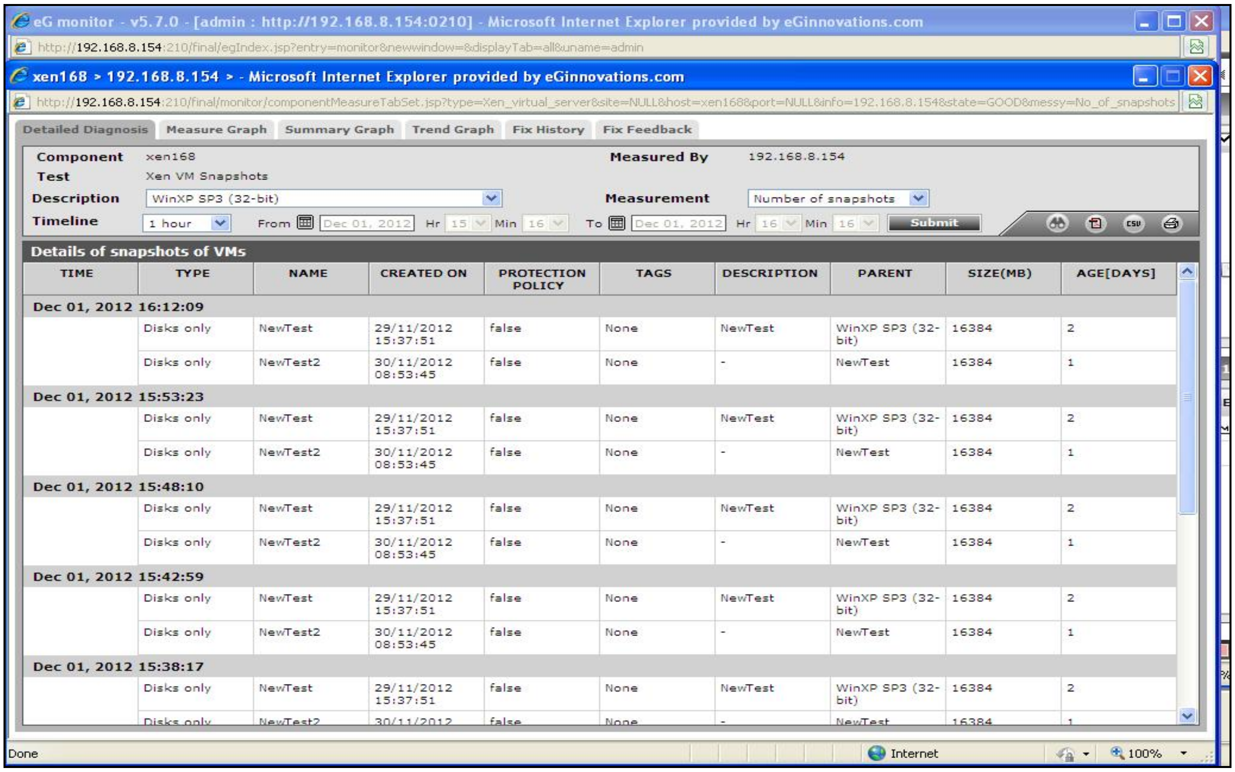
eG monitor - v5.7.0 - [admin : http://192.168.8.154:0210] - Microsoft Internet Explorer provided by eGinnovations.com
http://192.168.8.154:210/final/egIndex.jsp?entry=monitor&newwindow=&displayTab=all&uname=admin
xen168 > 192.168.8.154 > - Microsoft Internet Explorer provided by eGinnovations.com
http://192.168.8.154:210/final/monitor/componentMeasureTabSet.jsp?type=Xen_virtual_server&site=NULL&host=xen168&port=NULL&info=192.168.8.154&state=GOOD&messy=No_of_snapshots
Detailed Diagnosis
Measure Graph
Summary Graph
Trend Graph
Fix History
Fix Feedback
Submit
Done
Internet
100%

Component: xen168
Test: Xen VM Snapshots
Measured By: 192.168.8.154
Description: WinXP SP3 (32-bit)
Measurement: Number of snapshots
Timeline: 1 hour From: Dec 01, 2012 Hr 15 Min 16 To: Dec 01, 2012 Hr 16 Min 16

Details of snapshots of VMs

TIME	TYPE	NAME	CREATED ON	PROTECTION POLICY	TAGS	DESCRIPTION	PARENT	SIZE(MB)	AGE[DAYS]
Dec 01, 2012 16:12:09									
	Disks only	NewTest	29/11/2012 15:37:51	false	None	NewTest	WinXP SP3 (32-bit)	16384	2
	Disks only	NewTest2	30/11/2012 08:53:45	false	None	-	NewTest	16384	1
Dec 01, 2012 15:53:23									
	Disks only	NewTest	29/11/2012 15:37:51	false	None	NewTest	WinXP SP3 (32-bit)	16384	2
	Disks only	NewTest2	30/11/2012 08:53:45	false	None	-	NewTest	16384	1
Dec 01, 2012 15:48:10									
	Disks only	NewTest	29/11/2012 15:37:51	false	None	NewTest	WinXP SP3 (32-bit)	16384	2
	Disks only	NewTest2	30/11/2012 08:53:45	false	None	-	NewTest	16384	1
Dec 01, 2012 15:42:59									
	Disks only	NewTest	29/11/2012 15:37:51	false	None	NewTest	WinXP SP3 (32-bit)	16384	2
	Disks only	NewTest2	30/11/2012 08:53:45	false	None	-	NewTest	16384	1
Dec 01, 2012 15:38:17									
	Disks only	NewTest	29/11/2012 15:37:51	false	None	NewTest	WinXP SP3 (32-bit)	16384	2
	Disks only	NewTest2	30/11/2012	false	None	-	NewTest	16384	1

Figure 3.6: The detailed diagnosis of the Number of snapshots measure

3.1.5 Xen VM Tools Status Test

XenServer Tools, the XenServer paravirtualization tools, provide high speed I/O for enhanced disk and network performance. XenServer Tools must be installed on each virtual machine (Windows and Linux) in order to use the xe CLI or XenCenter; moreover, VM performance will significantly deteriorate unless the tools are installed.

Without the tools installed, you will not be able to do any the following:

- Cleanly shut down, reboot or suspend a VM.
- View VM performance data in XenCenter.
- Migrate a running VM (via XenMotion).
- Create quiesced snapshots or snapshots with memory (checkpoints), or revert to snapshots.
- Adjust the number of vCPUs on a running Linux VM (Windows VMs require a reboot for this to take effect).

Using the **Xen VM Tools Status** test, you can ascertain whether the XenServer tools are installed or not on a VM.

Target of the test : A XenServer host

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each VM on the target XenServer host

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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.
9. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.
10. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to *1:1*, it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Is Xen tools installed?:	Indicates whether/not XenServers Tools has been		If XenServer Tools is installed on a VM, this measure will report the value Yes.

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td>installed on this VM.</td><td></td><td>On the other hand, if XenServer Tools is not installed on a VM, this measure will report the value No.
The numeric values that correspond to each of the measure values discussed above have been detailed 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>
<b>Note:</b>
By default, this measure reports <i>Yes</i> or <i>No</i> to indicate whether XenServer Tools is installed on a VM or not. The graph of this measure however, represents the same using the numeric equivalents only – i.e., 0 or 1.</td></tr>
<tr><td><b>Is Xen tools upto date?:</b></td><td>Indicates whether the latest version of XenServer Tools has been installed on this VM.</td><td></td><td>If the latest version of XenServer Tools is installed on a VM, this measure will report the value Yes. On the other hand, if an older/obsolete version of XenServer Tools is installed on a VM, this measure will report the value No.
The numeric values that correspond to each of the measure values discussed above have been detailed 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>
<b>Note:</b>
By default, this measure reports <i>Yes</i> or <i>No</i> to indicate the version of XenServer</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			Tools installed on a VM. The graph of this measure however, represents the same using the numeric equivalents only – i.e., 0 or 1.

3.1.6 VDI Applications Test

This test discovers the applications executing on the virtual desktops and reports the availability and resource-usage of each of the desktop applications.

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of outputs for every distinct application executing on the virtual desktops

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right

to connect to the XenServer console via SSH.

- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **IS_SHOW_ALL_APPS** - To ensure that the test monitors only specific applications executing on the desktops and not all of them, set the **IS_SHOW_ALL_APPS** flag to **No**. Once this is done, then, you need to configure those applications that you want to exclude from the monitoring scope of this test. For this purpose, follow the steps given below:

    - Edit the **eg_tests.ini** file (in the **{EG_INSTALL_DIR}\manager\config directory**).
    - In the **[EXCLUDE_APPLICATIONS]** section of the file, you will find an entry of the following format:

      VmgApplicationTest={Comma-separated list of applications to be excluded}

    - To the comma-separated application list that pre-exists, append the applications that you want to monitor. For instance, if your test need not monitor notepad.exe, and powerpnt.exe, then, your entry should be:

      VmgApplicationTest=..................,notepad.exe,powerpnt.exe

      Note that the exact application names should be provided, but the extensions (for instance, .exe) can be dispensed with.

    - Finally, save the file.

    On the other hand, if you want to monitor all the applications, then, set the **IS_SHOW_ALL_APPS** flag to **Yes**.

12. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

13. **SHOW USER APPS ONLY** – By default, this flag is set to **Yes**. Accordingly, this test will monitor only those applications/processes that are running in the user's account. To monitor all applications/processes running in the virtual desktops, regardless of the user account using which they are running, set this flag to **No**.

14. **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.

15. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to *1:1*, it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.

16. **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.

Measurements of the test

Measurement	Description	Measurement Unit	Interpretation
Processes running:	Indicates the number of instances of this application that is currently executing across all virtual desktops on the target host operating system.	Number	
CPU usage:	Indicates the percentage of physical CPU resources utilized by this application across the guest VMs.	Percent	A very high value of this measure is a cause for concern, as it indicates excessive CPU usage by a single application. This in turn would cause other desktop applications to contend for limited physical resources, thus degrading the performance of those applications and that of the virtual server as a whole.
Memory usage:	Indicates the percentage of physical memory resources	Percent	A very high value of this measure is a cause for concern, as it indicates

Measurement	Description	Measurement Unit	Interpretation
	utilized by this application across the guest VMs.		excessive memory usage by a single application. This in turn would cause other desktop applications to contend for limited physical memory resources, thus degrading the performance of those applications and that of the virtual server as a whole.
CPU used:	Indicates the physical CPU (in Mhz) used up by this application.	Mhz	
I/O reads:	Indicates the rate at which this application reads data from the virtual disk.	KB/Sec	Compare the value of these measures across applications to identify the I/O-intensive application.
I/O writes:	Indicates the rate at which this application writes data to the virtual disk.	KB/Sec	
I/O read operations:	Indicates the rate at which this application performs read operations.	IOPS	Compare the value of these measures across applications to identify the I/O-intensive application.
I/O write operations:	Indicates the rate at which this application performs write operations.	IOPS	
Total I/O operations:	Indicates the rate at which this application performs I/O operations.	IOPS	

3.2 The Inside View of Desktops 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 XenServer host is 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 Desktops** layer provide an "internal" view of the workings of each of the desktops - these tests 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 user sessions, TCP traffic, and network loading.

By default however, clicking on the **Inside View of Desktops** layer, does not list the associated tests. Instead, Figure 3.7 appears, which displays the current state of all virtual desktops that have been configured on the monitored XenServer host.

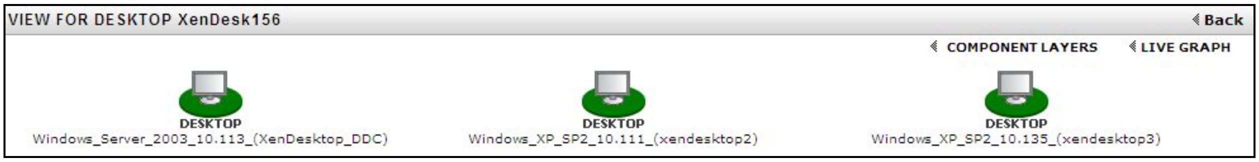
VIEW FOR DESKTOP XenDesk156
Back
COMPONENT LAYERS
LIVE GRAPH
DESKTOP
Windows_Server_2003_10.113_(XenDesktop_DDC)
DESKTOP
Windows_XP_SP2_10.111_(xendesktop2)
DESKTOP
Windows_XP_SP2_10.135_(xendesktop3)

Figure 3.7: The current state of the desktops configured on the XenServer host that is monitored

If you want to override this default setting - i.e., if you prefer to view the tests mapped to the **Inside View of Desktops** layer first, and then proceed to focus on individual desktop 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 Desktops** layer is clicked, the list of tests mapped to that layer appears.

From the desktop view, you can further drill-down to focus on the health of a particular desktop, by clicking on the icon representing the desktop in Figure 3.7. Figure 3.8 then appears displaying all the performance metrics extracted from that virtual desktop 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 desktop. To view the time-of-day variations in a measure, you can view its graph by clicking on that measure in Figure 3.8.

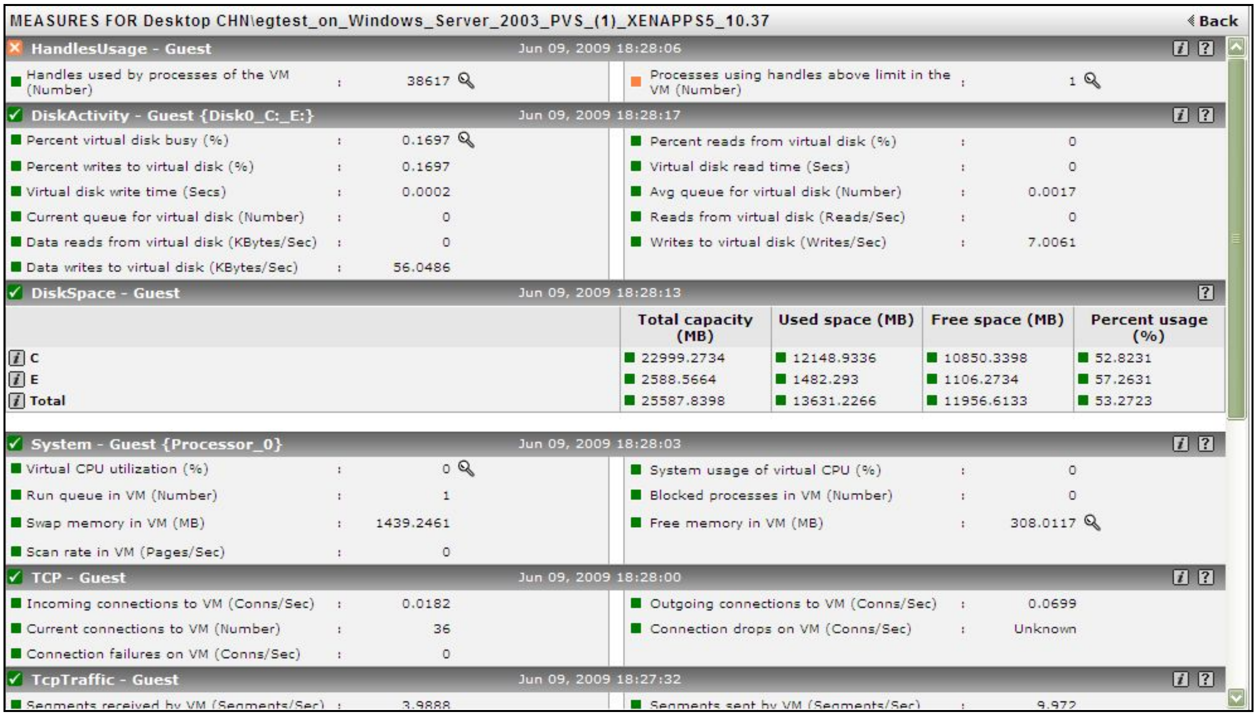

MEASURES FOR Desktop CHN\egtest_on_Windows_Server_2003_PVS_(1)_XENAPPS5_10.37 Back

☒ HandlesUsage - Guest Jun 09, 2009 18:28:06

- Handles used by processes of the VM (Number) : 38617
- Processes using handles above limit in the VM (Number) : 1

✓ DiskActivity - Guest {Disk0_C:_E:} Jun 09, 2009 18:28:17

- Percent virtual disk busy (%) : 0.1697
- Percent writes to virtual disk (%) : 0.1697
- Virtual disk write time (Secs) : 0.0002
- Current queue for virtual disk (Number) : 0
- Data reads from virtual disk (KBytes/Sec) : 0
- Data writes to virtual disk (KBytes/Sec) : 56.0486
- Percent reads from virtual disk (%) : 0
- Virtual disk read time (Secs) : 0
- Avg queue for virtual disk (Number) : 0.0017
- Reads from virtual disk (Reads/Sec) : 0
- Writes to virtual disk (Writes/Sec) : 7.0061

✓ DiskSpace - Guest Jun 09, 2009 18:28:13

	Total capacity (MB)	Used space (MB)	Free space (MB)	Percent usage (%)
C	22999.2734	12148.9336	10850.3398	52.8231
E	2588.5664	1482.293	1106.2734	57.2631
Total	25587.8398	13631.2266	11956.6133	53.2723

✓ System - Guest {Processor_0} Jun 09, 2009 18:28:03

- Virtual CPU utilization (%) : 0
- Run queue in VM (Number) : 1
- Swap memory in VM (MB) : 1439.2461
- Scan rate in VM (Pages/Sec) : 0
- System usage of virtual CPU (%) : 0
- Blocked processes in VM (Number) : 0
- Free memory in VM (MB) : 308.0117

✓ TCP - Guest Jun 09, 2009 18:28:00

- Incoming connections to VM (Conns/Sec) : 0.0182
- Current connections to VM (Number) : 36
- Connection failures on VM (Conns/Sec) : 0
- Outgoing connections to VM (Conns/Sec) : 0.0699
- Connection drops on VM (Conns/Sec) : Unknown

✓ TcpTraffic - Guest Jun 09, 2009 18:27:32

- Segments received by VM (Segments/Sec) : 3.9888
- Segments sent by VM (Segments/Sec) : 9.972

Figure 3.8: The measures pertaining to a particular desktop

You can also view live graphs of pre-configured measures pertaining to a XenServer host and the virtual desktops configured on it, by clicking on the **LIVE GRAPHS** link in Figure 3.7. Alternatively, you can click on the icon that appears in the **Tests** panel when the **Inside View of Desktops** layer is clicked in the layer model page to view the live graph. The resulting graph display (see Figure 3.9) has been organized in such a way that next to every host-pertinent measure graph, the closely related guest-specific measure graph appears. For instance, next to the graph of the ‘Physical CPU usage’ measure of the XenCpu test, you will find a graph of the 'Physical CPU usage' measure of the XenxGuestDetails test. This way, you can easily compare and correlate how well the physical CPU resources are being utilized by both the XenServer host 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 host? or is it the virtual guest?

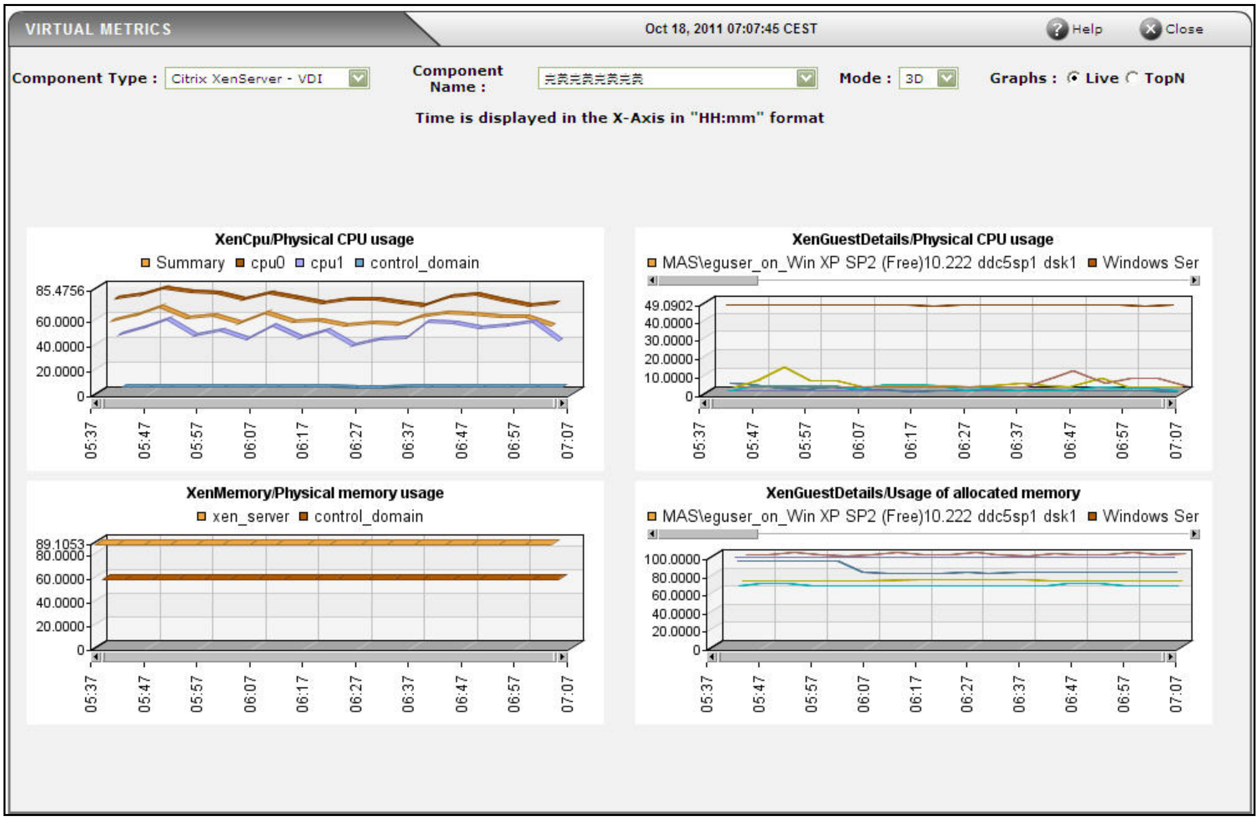

VIRTUAL METRICS

Oct 18, 2011 07:07:45 CEST

Component Type : Citrix XenServer - VDI

Component Name : 完莞完莞完莞完莞

Mode : 3D

Graphs : ● Live ○ TopN

Time is displayed in the X-Axis in "HH:mm" format

Figure 3.9: Live graph for Citrix XenServer - VDI

To return to the layer model of the *Citrix XenServer - VDI* server and view the tests mapped to the **Inside View of Desktops** layer, click on the **COMPONENT LAYERS** link in Figure 3.7. The tests depicted by Figure 3.10 then appear.

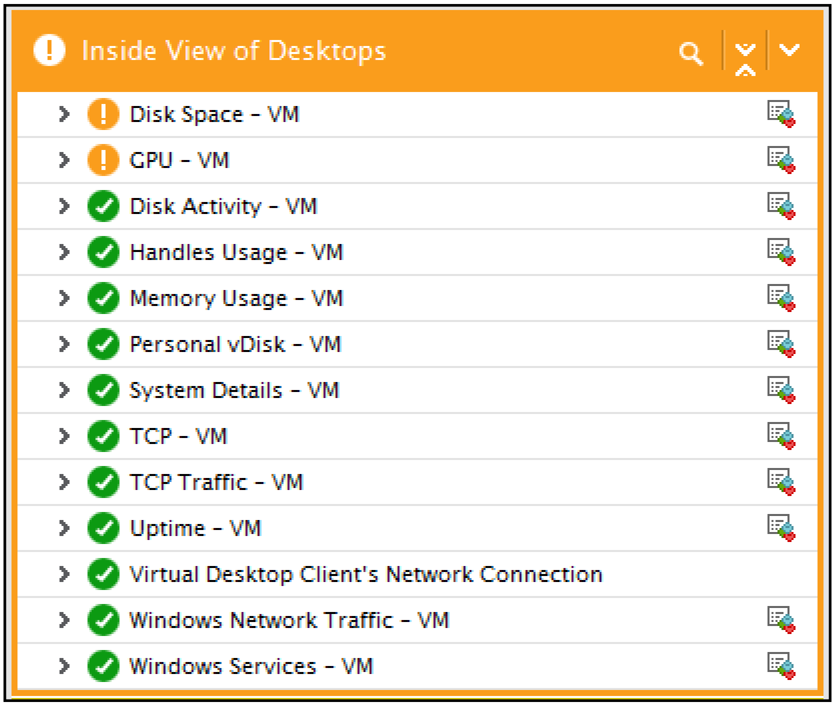

Inside View of Desktops

- Disk Space - VM
- GPU - VM
- Disk Activity - VM
- Handles Usage - VM
- Memory Usage - VM
- Personal vDisk - VM
- System Details - VM
- TCP - VM
- TCP Traffic - VM
- Uptime - VM
- Virtual Desktop Client's Network Connection
- Windows Network Traffic - VM
- Windows Services - VM

Figure 3.10: The tests associated with the Inside View of Desktops layer of a Citrix XenServer - VDI

Almost all the tests depicted by Figure 3.10 have already been dealt with in the Section chapter. The tests that are specific to the *Citrix XenServer - VDI* model are the Terminal to Desktop Connection test, Desktop ICA Channel Test, and the PCoIP Session - VM test.

3.2.1 Terminal to Desktop Connection Test

A Virtual Desktop Infrastructure (VDI) is a shared environment in which multiple users connect to desktops hosted by virtual machines executing on a XenServer host from remote terminals. One of the key factors influencing user experience in such an environment is the latency seen by the users when connecting to a virtual desktop. High network latencies or packet losses during transmission can cause significant slow-downs in request processing by the desktop. Hence, monitoring latencies between the virtual desktop and individual client terminals is important.

The Terminal to Desktop Connection test is executed by the eG agent on a XenServer host. This test auto-discovers the virtual desktops on the XenServer, the users who are currently logged on to each of the virtual desktops, and the IP address from which they are connecting to the virtual desktops. For each user, the test monitors the quality of the link between the client and the virtual desktop.

Using this test, an administrator can identify user sessions that are being impacted by high latencies or by excessive packet drops. In some cases, a virtual desktop may regard a user session as active, even though the network link connecting the user terminal to the virtual desktop has failed. The Terminal to Desktop Connection test alerts administrators to such situations.

Note:

This test will work only on Windows VMs.

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of outputs for every user currently connected to the virtual desktop

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

- Obtain the server-certificate for the XenServer
- Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-

enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

14. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

15. **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 *virtualmachine* 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.

16. **PACKETSIZE** - The size of packets used for the test (in bytes)

17. **PACKETCOUNT** - The number of packets exchanged between the virtual desktop and the user terminal during the test

18. **TIMEOUT** - How long after transmission should a packet be deemed lost (in seconds)

19. **PACKETINTERVAL** - Represents the interval (in milliseconds) between successive packet transmissions during the execution of this test.

20. **REPORTUNAVAILABILITY** – By default, this flag is set to No. This implies that, by default, the test will not report the unavailability of network connection between a user terminal and a virtual desktop. In other words, if the Packet loss measure of this test registers the value 100% for any user, then, by default, this test will not report any measure for that user; under such circumstances, the corresponding user name will not appear as a descriptor of this test. You can set this flag to Yes, if you want the test to report and alert you to the unavailability of network connection between a user terminal and a virtual desktop.

Measurements of the test

Measurement	Description	Measurement Unit	Interpretation
Number of	Indicates the current	Number	The value 0 indicates that the user is not

Measurement	Description	Measurement Unit	Interpretation
sessions:	number of sessions for a particular user		currently connected to the virtual desktop.
Average delay:	Indicates the average delay between transmission of a request by the agent on a virtual desktop and receipt of the response back from the user terminal.	Secs	Comparing the value of this measure across users will enable administrators to quickly and accurately identify users who are experiencing higher latency when connecting to a virtual desktop.
Minimum delay:	Indicates the minimum delay between transmission of a request by the agent on a virtual desktop and receipt of the response back from the user terminal.	Secs	A significant increase in the minimum round-trip time is often a sure sign of a poor link between the desktop and a user's terminal.
Packet loss:	Indicates the percentage of packets lost during data exchange between the virtual desktop and the user terminal.	Percent	Comparing the value of this measure across users will enable administrators to quickly and accurately identify users who are experiencing slowdowns because of poor performance on the network links between their terminals and the virtual desktop.

Note:

- If the same user is connecting to the virtual desktop from multiple client terminals, the value of the *Number of sessions*, *Average delay*, and *Packet loss* measures will be averaged across all the sessions of that user. The *Minimum delay* measure, on the other hand, will display the least value reported for *Minimum delay* across all the sessions of that user.
- When a user logs out, the number of sessions will be reduced by 1. If the number of user sessions becomes 0, the corresponding entry for that user in the eG user interface will be removed after a short period of time.

3.2.2 Desktop’s HDX Channel Test

As already mentioned, the key factors influencing user experience in a virtual desktop infrastructure are the latencies experienced by the user while connecting to the desktop via ICA and the bandwidth used when a user interacts with a virtual desktop. High latency and excessive bandwidth consumption can often slowdown access to desktops, thereby significantly delaying subsequent user operations. Hence, monitoring the latency and bandwidth usage of the ICA communication channel between the user terminal and the virtual desktops is essential.

The Desktop's HDX Channel test auto-discovers the virtual desktops on the XenServer host and the users who are currently connected to each desktop. For each such user, the test monitors the communication between a user and the virtual desktop, and reports the following:

- The latency experienced by each user session;
- The bandwidth used by the incoming and outgoing data/audio/multimedia traffic transacted by the ICA communication channel between each user and virtual desktop;

Using this test, an administrator can identify user sessions that are being impacted by high latency and abnormal bandwidth usage. In addition, the test also reveals the type of traffic that is causing excessive bandwidth usage, thereby providing pointers to how the client configuration can be fine-tuned in order to reduce bandwidth consumption and improve performance.

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

Note:

This test will report metrics only if the following conditions are fulfilled:

- The test is applicable to **Windows VMs** only.
- The VMs being monitored should be managed by XenDesktop Broker.
- The Virtual Desktop Agent software should have been installed on the VMs.
- The **ICA Session** performance object should be enabled on the VMs.

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every user who is connected to a virtual desktop via ICA

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can

configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

- Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
- In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
- Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

   - **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 PASSWORD**s 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 **2.5.1.1**.

   - **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

14. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

15. **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 *virtualmachine* 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.

Measurements of the test

Measurement	Description	Measurement Unit	Interpretation
Session average latency:	Indicates the average client latency over the lifetime of this session.	Secs	Comparing the value of this measure across users will enable administrators to quickly and accurately identify users who are experiencing higher latency when connecting to a virtual desktop.
Session deviation latency:	Indicates the difference between the minimum and maximum measured latency values for this session.	Secs	
Audio bandwidth input:	Indicates the bandwidth used while transmitting sound/audio to this user.	Kbps	Comparing these values across users will reveal which user is sending/receiving bandwidth- intensive sound/audio files over the ICA channel. To minimize bandwidth consumption, you may want to consider disabling client audio mapping.
Audio bandwidth output:	Indicates the bandwidth used while receiving sound/audio from this user.	Kbps	
COM bandwidth input:	Indicates the bandwidth used when sending data to this user's COM port.	Kbps	Comparing these values across users will reveal which user's COM port is sending/receiving bandwidth- intensive data over the ICA channel.
COM bandwidth ouput:	Indicates the bandwidth used when receiving data from this user's COM port.	Kbps	
Drive bandwidth	Indicates the bandwidth	Kbps	Comparing the values of these measures

Measurement	Description	Measurement Unit	Interpretation
input:	used when this user performs file operations on the mapped drive on the virtual desktop.		across users will reveal which user is performing bandwidth- intensive file operations over the ICA channel. If bandwidth consumption is too high, you may want to consider disabling client drive mapping on the client device. Client drive mapping allows users logged on to a virtual desktop from a client device to access their local drives transparently from the ICA session. Alternatively, you can conserve bandwidth by even refraining from accessing large files with client drive mapping over the ICA connection.
Drive bandwidth output:	Indicates the bandwidth used when the virtual desktop performs file operations on the client's drive.	Kbps	
Printer bandwidth input:	Indicates the bandwidth used when this user prints to a desktop printer over the ICA channel.	Kbps	Comparing the values of these measures across users will reveal which user is issuing bandwidth- intensive print commands over the ICA channel. If bandwidth consumption is too high, you may want to consider disabling printing. Alternatively, you can avoid printing large documents over the ICA connection.
Printer bandwidth output:	Indicates the bandwidth used when the desktop responds to print jobs issued by this user.	Kbps	
Session bandwidth input:	Indicates the bandwidth used from this user to the virtual desktop for a session	Kbps	Comparing the values of these measures across users will reveal which user and which virtual desktop is performing bandwidth- intensive operatons for a session.
Session bandwidth output:	Indicates the bandwidth used from the virtual desktop to this user for a session.	Kbps	
Session compression input:	Indicates the compression ratio used from this user to the virtual desktop for a session.	Number	Compression reduces the size of the data that is transacted over the ICA channel.

Measurement	Description	Measurement Unit	Interpretation
			Comparing the values of these measures across users will reveal which client has been configured with a very low and a very high compression ratio. In the event of high bandwidth usage over an ICA channel, you can set a higher compression ratio for the corresponding client and thus reduce bandwidth consumption.
Session compression output:	Indicates the compression ratio used from the virtual desktop to this user for a session.	Number	
Speed screen data channel bandwidth input:	Indicates the bandwidth used from this user to the virtual desktop for data channel traffic.	Kbps	Comparing the values of these measures across users will reveal which user has been transmitting/receiving bandwidth-intensive data channel traffic.
Speed screen data channel bandwidth output:	Indicates the bandwidth used from virtual desktop to this user for data channel traffic.	Kbps	
Speed screen multimedia acceleration bandwidth input:	Indicates the bandwidth used from this user to virtual desktop for multimedia traffic.	Kbps	Comparing the values of these measures across users will reveal which user has been transmitting/receiving bandwidth-intensive multimedia traffic.
Speed screen multimedia acceleration bandwidth output:	Indicates the bandwidth used from the virtual desktop to this user for multimedia traffic	Kbps	
HDX media stream for flash data bandwidth input:	Indicates the bandwidth used from this user to virtual desktop for flash data traffic.	Kbps	Comparing the values of these measures across users will reveal which user has been transmitting/receiving bandwidth-intensive flash data.
HDX media stream	Indicates the bandwidth	Kbps	

Measurement	Description	Measurement Unit	Interpretation
for flash data bandwidth output:	used from the virtual desktop to this user for flash data traffic		
USB bandwidth input:	Indicates the bandwidth used from this user to the virtual desktop for the USB port-related traffic.	Kbps	Comparing the values of these measures across users will reveal which user has been transmitting/receiving bandwidth-intensive USB traffic.
USB bandwidth output:	Indicates the bandwidth used from the virtual desktop to this user for the USB port-related traffic.	Kbps	
Last recorded latency:	Indicates the last recorded latency of this user session.	Secs	Comparing the value of this measure across user sessions will enable administrators to quickly and accurately identify users who experienced high latencies recently.
Input line speed:	Indicates the average line speed of all the sessions of this user to the desktop.	KB/Sec	
Output line speed:	Indicates the average line speed of all sessions from the desktop to this user.	KB/Sec	
Bandwidth usage:	Indicates the percentage HDX bandwidth consumption of this user.	Percent	Compare the value of this measure across users to know which user is consuming the maximum HDX bandwidth.
Resource shares:	Indicates the total number of resource shares used by this user.	Number	By comparing the value of this measure across users, you can identify the user who is hogging the resources.
Frame rate:	Indicates the rate at which frames are processed	Frames/Sec	FPS is how fast your graphics card can output individual frames each second. It

Measurement	Description	Measurement Unit	Interpretation
	during this user session.		is the most time- tested and ideal measure of performance of a GPU. Higher the value of this measure, healthier is the GPU.
Framehawk frame rate:	Indicates the rate at which frames are processed by the Framehawk virtual channel, if it is enabled for this user session.	Frames/Sec	The Framehawk virtual channel optimizes the delivery of virtual desktops and applications to users on broadband wireless connections, when high packet loss or congestion occurs. **Note:** This measure will report the value 0 if Framehawk is not enabled for a user or if the device from which the user is accessing the application does not support Framehawk.
Framehawk network bandwidth:	Indicates the bandwidth consumption of this user session when the Framehawk virtual delivery channel is used.	Kbps	This measure will report the value 0 if Framehawk is not enabled for a user or if the device from which the user is accessing the application does not support Framehawk.
Framehawk latency:	Indicates the latency experienced by this user session when the Framehawk virtual delivery channel is used.	Secs	This measure will report the value 0 if Framehawk is not enabled for a user or if the device from which the user is accessing the application does not support Framehawk.
Framehawk network loss:	Indicates the percentage of packet loss experienced by this user session when the Framehawk virtual delivery channel is used.	Percent	This measure will report the value 0 if Framehawk is not enabled for a user or if the device from which the user is accessing the application does not support Framehawk.

3.2.3 PCoIP Session - VM Test

PCoIP - PC over IP - is a proprietary protocol for remote workstation and desktop resolution. VMware View supports PCoIP to deliver virtual desktops to users connecting to the VDI. Since PCoIP recognizes different types of content and then uses different compression algorithms based on the content type, it is often considered ideal to deliver on the VDI promise of a rich user experience.

The key factors influencing user experience in such cases are the latencies experienced by the user while connecting to the desktop via PCoIP and the bandwidth used when a user interacts with a virtual desktop. High latency and excessive bandwidth consumption can often slowdown access to desktops, thereby significantly delaying subsequent user operations. Hence, monitoring the latency and bandwidth usage of the PCoIP communication channel between the user terminal and the virtual desktops is essential.

The PCoIP Session - VM test auto-discovers the virtual desktops on the XenServer host and the users who are currently connected to each desktop via PCoIP. For each such user, the test monitors the communication between a user and the virtual desktop, and reports the following:

- The latency experienced by each user session;
- The bandwidth used by the incoming and outgoing data/audio/multimedia traffic transacted by the PCoIP communication channel between each user and virtual desktop;

Using this test, an administrator can identify user sessions that are being impacted by high latency and abnormal bandwidth usage. In addition, the test also reveals the type of traffic that is causing excessive bandwidth usage, thereby providing pointers to how the client configuration can be fine-tuned in order to reduce bandwidth consumption and improve performance.

This test is relevant only where VMware View is used to broker connections between the user and the desktops. Hence, this test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence: Agents -> Tests -> Enable/Disable, pick *Citrix XenServer - VDI* as the **Component type**, set *Performance* as the **Test type**, choose this test from the **DISABLED TESTS** list, and click on the **>>** button to move the test to the **ENABLED TESTS** list.

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every user who is connected to a virtual desktop via PCoIP

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the

credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

- Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
- In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
- Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

   - **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 PASSWORD**s 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

2.5.1.1.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

14. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

15. **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 *virtualmachine* 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.

Measurements of the test

Measurement	Description	Measurement Unit	Interpretation
Round trip time:	Indicates the round trip latency between the virtual desktop and this user terminal.	Secs	Comparing the value of this measure across users will enable administrators to quickly and accurately identify users who are experiencing higher latency when connecting to a virtual desktop.
Data received rate:	Indicates the rate at which data was received by this user from the virtual desktop.	Kbit/Sec	Comparing the value of each of these measures across users will enable administrators to quickly and accurately identify users who are consuming the maximum bandwidth. Once you zero-in on the user, you can compare the Data received rate of that user with the Data sent rate to know when the user consumed more bandwidth - when receiving data or while sending data?
Data sent rate:	Indicates the rate at which data was sent by this user to the virtual desktop.	Kbit/Sec	
Audio data received:	Indicates the bandwidth used while transmitting sound/audio to this user.	Kbit/Sec	Comparing these values across users will reveal which user is sending/receiving bandwidth- intensive sound/audio files over PCoIP.
Audio data sent:	Indicates the bandwidth used while receiving sound/audio from this user.	Kbit/Sec	
Imaging data received rate:	Indicates the bandwidth used when sending imaging data to this user.	Kbit/Sec	Comparing these values across users will reveal which user is sending/receiving bandwidth- intensive images over PCoIP.

Measurement	Description	Measurement Unit	Interpretation
Imaging data sent rate:	Indicates the bandwidth used when receiving imaging data from this user.	Kbit/Sec	
Imaging decoder capability rate:	Indicates the currrent estimate of the decoder processing capability.	Kbit/Sec	
Incoming bandwith:	Indicates the overall bandwidth used by incoming PCoIP packets.	Kbit/Sec	Comparing the values of these measures across users will reveal which user is performing bandwidth-intensive operations over the PCoIP channel.
Outgoing bandwidth:	Indicates the overall bandwidth used by outgoing PCoIP packets.	Kbit/Sec	
USB data received rate:	Indicates the bandwidth used when this user received USB data over the PCoIP channel.	Kbit/Sec	Comparing the values of these measures across users will reveal which user is sending/receiving bandwidth-intensive USB data over the PCoIP channel.
USB data sent rate:	Indicates the bandwidth used when this user sent USB data over the PCoIP channel.	Kbit/Sec	
Received packets lost:	Indicates the percentage of packets received by this user that were lost.	Percent	A high value for these measures is indicative of a bad network connection between the user terminal and the virtual desktop.
Transmitted packets lost:	Indicates the percentage of packets transmitted by this user that were lost.	Percent	
Imaging encoded frames:	Indicates the number of imaging frames that were encoded per second.	Frames/Sec	

3.2.4 User Profile Management - Guest

User logon is a complex and resource intensive process in a VDI environment, and is a key determinant of the quality of a user's experience with the VDI service. This process is initiated when a desktop broker's load balancing algorithm selects the virtual desktop where a published application or desktop, which a user has selected, will be started and ends when the application or desktop is running and the user is able to interact with it.

Delays in the user logon process can therefore serve as key spoilers of a user's experience with the desktop service, causing significant loss of revenue and reputation in mission-critical VDI environments.

One of the common causes for delays in user logons is a delay in the loading of user profiles. To reduce the time taken to load profiles and thus minimize the user logon time, VDI environments where user connections are brokered through the Citrix XenDesktop Broker, use the Citrix Profile Management solution. Citrix Profile Management is a profile type that supersedes all other profiles for the user.

During logon, the Profile management service manages the user settings in a user profile. This service helps minimize the user logon time by enabling administrators to exclude (and include) certain files and folders in order to prevent extraneous settings from needlessly being copied with the profile. For example, some applications may create folders and files that account for tens or hundreds of megabytes—data that is really not required. By excluding these items, the profile is thus smaller, and smaller profiles load faster. Alternatively, you could elect to only include specific files and folders, thus keeping to a minimum the amount of profile data being managed within the user‘s profile.

Also, upon logoff, the Profile management service merges back only changed user settings to the centrally stored user settings (user‘s store).

In environments where the Citrix Profile Management service is utilized therefore, the user experience with the VDI service greatly depends upon how efficient the service is.

To ascertain the efficiency of the Citrix Profile Management service, VDI administrators may have to periodically track the logon/logoff duration and profile size of each user to the virtual desktops operating on a target virtual host. Doing so will enable these administrators to determine whether/not the Profile management service has succeeded in minimizing both user logon times and profile sizes. The **User Profile Management - Guest** test helps administrators perform this check at pre-configured intervals. The 'per-user' performance results reported by this test will not only enable administrators to judge the effectiveness of the Profile management service in its entirety, but will also shed light on those user logons/logoffs that are still experiencing delays; this provides insights into how the service can be fine-tuned to enhance the VDI experience of such users.

Note:

This test is relevant only where the Citrix XenDesktop Broker is used to broker connections between the user and the desktops

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

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every user to each virtual desktop on the Citrix XenServer being monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use

by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

- Obtain the server-certificate for the XenServer
- Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

14. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

15. **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 *virtualmachine* 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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Logon Duration:	Indicates the duration of logon processing for this user.	Secs	This value helps to measure the reduction in logon times when the Profile Management service 'streams' the profile. Ideally therefore, this value should be low. A high value or a consistent increase in the value of this measure could indicate that profile loading still takes a lot of time at logon - this could be owing to a large profile size. You can then check the value reported by the Logon Bytes measure to know the profile size at logon. If profile sizes continue to grow at logon despite the use of Profile management, it is

Measurement	Description	Measurement Unit	Interpretation
			indicative of the ineffectiveness of profile management. You may then have to fine-tune the feature to further reduce the profile size by excluding more unnecessary files from the profile, or you may have to explore other options such as roaming profiles, mandatory profiles, etc.
Logon Bytes:	Indicates the size of this user's profile when it is retrieved from the user's store at logon.	MB	Ideally, the value of this measure should be low. A low profile size could result in faster profile loading at logon, below time to login, and consequently, a richer user experience with the VDI service. If profile sizes continue to grow despite the use of Profile management, it is indicative of the ineffectiveness of profile management. You may then have to fine-tune the feature to further reduce the profile size by excluding more unnecessary files from the profile.
Logoff Duration:	Indicates the duration of logoff processing for this user.	Secs	A low value is desired for this measure. A high value could indicate that the profile management service takes too long to update the user's store with changes in the user settings. This could be because of a bad network connection between the virtual desktop and the user's store, or because too many changes are waiting to be written to the user store.
Logoff Bytes:	Indicates the size of this user's profile when it is copied to the user store at logoff.	MB	This measure provides a fair idea of the volume of changes that were copied to the user's store at logoff.
Local Profile Setup	Indicates the time taken to	Secs	A low value is desired for these

Measurement	Description	Measurement Unit	Interpretation
Duration:	create or prepare this user's profile on the local computer.		measures. If a user complaints of delays during logon, you can use the value of these measures to determine where the VDI service is spending too much time - is it when setting up the local profile? or is it when deleting the local profile?
Delete Local Profile Duration:	Indicates the time spent deleting this user's local profiles during the initial migration.	Secs	
Processed Logon Files - Under 1KB:	Indicates the number of locally copied files for this user's profile that are synchronized during logon and categorized by the file size of 1KB.	Number	All the **Processed Logon Files** measures help VDI administrators to understand whether/not 'profile streaming' (performed by the Profile Management service) has helped in reducing the number of locally copied files during logon. All the **Processed Logoff Files** measures help VDI administrators to understand how many files changed when the user session was in progress.
Processed Logoff Files - Under 1KB:	Indicates the number of locally copied file for this user's profile that are synchronized during logoff and categorized by the file size of 1KB.	Number	
Processed Logon Files - 1KB to 10KB:	Indicates the number of locally copied files for this user's profile that are synchronized during logon and categorized by the file size ranging from 1KB to 10KB.	Number	
Processed Logoff Files - 1KB to 10KB:	Indicates the number of locally copied files for this user's profile that are synchronized during logoff and categorized by the file size ranging from 1KB to 10KB.	Number	

Measurement	Description	Measurement Unit	Interpretation
Processed Logon Files - 10KB to 100KB:	Indicates the number of locally copied files for this user's profile that are synchronized during logon and categorized by the file size ranging from 10KB to 100KB.	Number	
Processed Logoff Files - 10KB to 100KB:	Indicates the number of locally copied files for this user's profile that are synchronized during logoff and categorized by the file size ranging from 1KB to 10KB.	Number	All the **Processed Logon Files** measures help VDI administrators to understand whether/not 'profile streaming' (performed by the Profile Management service) has helped in reducing the number of locally copied files during logon. All the **Processed Logoff Files** measures help VDI administrators to understand how many files changed when the user session was in progress.
Processed Logon Files - 100KB to 1MB:	Indicates the number of locally copied files for this user's profile that are synchronized during logon and categorized by the file size ranging from 100KB to 1MB.	Number	
Processed Logoff Files - 100KB to 1MB:	Indicates the number of locally copied files for this user's profile that are synchronized during logoff and categorized by the file size ranging from 100KB to 1MB.	Number	
Processed Logon Files - 1MB to 5MB:	Indicates the number of locally copied files for this user's profile that are synchronized during logon and categorized by the file size ranging from 1MB to 5MB.	Number	

Measurement	Description	Measurement Unit	Interpretation
Processed Logoff Files - 1MB to 5MB:	Indicates the number of locally copied files for this user's profile that are synchronized during logoff and categorized by the file size ranging from 1MB to 5MB.	Number	All the **Processed Logon Files** measures help VDI administrators to understand whether/not 'profile streaming' (performed by the Profile Management service) has helped in reducing the number of locally copied files during logon. All the **Processed Logoff Files** measures help VDI administrators to understand how many files changed when the user session was in progress.
Processed Logon Files - Above 5MB:	Indicates the number of locally copied files for this user's profile that are synchronized during logon and categorized by the file size above 5MB.	Number	
Processed Logoff Files - Above 5MB:	Indicates the number of locally copied files for this user's profile that are synchronized during logoff and categorized by the file size above 5MB.	Number	

3.2.5 Domain Time Sync – VM Test

Time synchronization is one of the most important dependencies of windows. A time protocol is responsible for determining the best available time information and converging the clocks to ensure that a consistent time is maintained across systems. By default, windows support a tolerance of plus or minus five minutes for clocks. If the time variance exceeds this setting, clients will be unable to authenticate and in the case of domain controllers, replication will not occur. It implements a time synchronization system based on Network Time Protocol (NTP).

NTP is a fault-tolerant, highly scalable time protocol and it is used for synchronizing computer clocks by using a designated reference clock. A reference clock is some device or machinery that spits out the current time. The special thing about these things is accuracy. Reference clocks must be accurately following some time standard. NTP will compute some additional statistical values based on the current time reported by the reference clock, which will describe the quality of time it sees. Among these values are: offset (or phase), jitter (or dispersion), frequency error, and stability. Thus each NTP server will maintain an estimate of the quality of its reference clocks and of itself.

This test reports the time difference between the reference clock and that of the target environment, and thus helps assess the quality of time seen by the Windows VM. With the help of this test, you can also easily determine whether the reference time changed recently.

Note:

This test reports metrics for Windows VMs only.

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

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each powered-on Windows guest/each user logged into a Windows guest on the Citrix XenServer monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.

6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.

7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

14. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

15. **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 *virtualmachine* 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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
NTP offset:	Indicates the time difference between the local clock and the designated reference clock.	Secs	For a tiny offset, NTP will adjust the local clock; for small and larger offsets, NTP will reject the reference time for a

Measurement	Description	Measurement Unit	Interpretation
			while. In the latter case, the operating system's clock will continue with the last corrections effective while the new reference time is being rejected. After some time, small offsets (significantly less than a second) will be slewed (adjusted slowly), while larger offsets will cause the clock to be stepped (set anew). Huge offsets are rejected, and NTP will terminate itself, believing something very strange must have happened.

3.2.6 Browser Activity – VM Test

When a user complains of a virtual desktop slowdown, administrators will have to instantly figure out if that VM is experiencing a resource crunch, and if so, which process/application on the desktop is contributing to it. One of the common reasons for CPU/memory contentions and handle leaks on a virtual desktop is web browsing! If a user to a virtual desktop browses resource-intensive web sites, it is bound to result in over-usage of the resources allocated to that VM, which in turn degrades the performance of not just that VM but even the other VMs on that host. While the **System Details – VM** test can lead administrators to the exact browser application that is consuming the CPU/memory resources of the VM excessively, it does not provide visibility into the precise websites that were been browsed when the resource contention occurred. This is where the **Browser Activity – VM** test helps.

For each web browser that is being accessed by a user per virtual desktop, this test reports how every browser uses the allocated CPU, memory, and disk resources and reveals the number and URLs of the web sites that are being accessed using each browser. This way, the test not only points administrators to resource-hungry browsers, but also indicates which web sites were being accessed using that browser.

Note:

- This test will report metrics only if the Windows VM being monitored uses the .Net framework v3.0 (or above).
- This test will not be able to monitor the Microsoft Edge browser on Windows 10 VMs.

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each browser used by every powered-on Windows guest/each user logged into a Windows guest on the Citrix XenServer monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.
5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can

specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

14. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

15. **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 *virtualmachine* 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.

16. **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**.

17. **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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Running browser instances:	Indicates the number of instances of this browser currently running on this	Number	Use the detailed diagnosis of this measure to know how much resources were utilized by each instance of a

Measurement	Description	Measurement Unit	Interpretation
	virtual desktop.		browser, so that the resource-hungry instance can be isolated.
Recent web sites:	Indicates the number of websites that were accessed using this browser on this virtual desktop during the last measurement period.	Number	Use the detailed diagnosis of this measure to know which web sites are being accessed using a browser.
CPU utilization:	Indicates the percentage CPU usage of this browser on this virtual desktop.	Percent	Compare the value of this measure across browsers to know which browser consumed the maximum CPU on a desktop. If the value of this measure is close to 100% on that desktop, it indicates excessive CPU usage by the browser. You may then want to use the detailed diagnosis of the Recent web sites measure to know which web sites are being accessed using that browser, which caused CPU usage to soar.
Memory used:	Indicates the percent usage of memory by this browser on this virtual desktop.	Percent	Compare the value of this measure across browsers to know which browser consumed the maximum memory on a desktop. If the value of this measure is close to 100% on that desktop, it indicates excessive memory usage by the browser. You may then want to use the detailed diagnosis of the Recent web sites measure to know which web sites are being accessed using that browser, which caused CPU usage to soar.
Handles used:	Indicates the number of handles opened by this browser on this virtual desktop.	Number	Compare the value of this measure across browsers to know which browser opened the maximum number

Measurement	Description	Measurement Unit	Interpretation
			of handles on a desktop. If the value of this measure consistently increases on that desktop, it indicates that the corresponding browser is leaking memory. You may then want to use the detailed diagnosis of the Recent web sites measure to know which web sites are being accessed using that browser, which caused the memory leak.
Disk reads:	Indicates the rate at which this browser read from the disks supported by this virtual desktop.	KB/Sec	A high value for these measures indicates that the browser is generating high disk I/O. You may then want to use the detailed diagnosis of the Recent web sites measure of this browser to know which web sites on the browser are responsible for the high disk I/O.
Disk writes:	Indicates the rate at which this browser read from the disks of this virtual desktop.	KB/Sec	
Disk IOPS:	Indicates the rate of read and write operations performed by this browser on the disks of this virtual desktop.	Operations/Sec	A high value for this measure indicates that the browser is generating high disk I/O. You may then want to use the detailed diagnosis of the Recent web sites measure of this browser to know which web sites on the browser are responsible for the high disk I/O.
Page faults:	Indicates the rate at which page faults by the threads executing in this browser are occurring on this virtual desktop.	Faults/Sec	Ideally, the value of this measure should be low. A high value for a browser is a cause for concern. You may then want to use the detailed diagnosis of the Recent web sites measure of this browser to know which web sites on the browser are responsible for page faults.

The detailed diagnosis of the **Running browser instances** measure reveals the process ID of each browser instance that is currently running on the virtual desktop and the resource usage of each instance. This way, you can easily and accurately identify the instance that is consuming resources excessively.

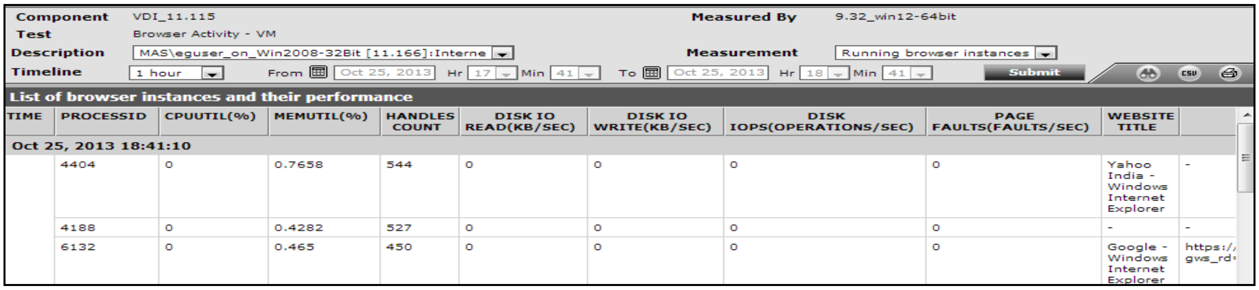

Component VDI_11.115
Test Browser Activity - VM
Description MAS\eguser_on_Win2008-32Bit [11.166]:Interne
Timeline 1 hour From Oct 25, 2013 Hr 17 Min 41 To Oct 25, 2013 Hr 18 Min 41
Measured By 9.32_win12-64bit
Measurement Running browser instances

List of browser instances and their performance

TIME	PROCESSID	CPUUTIL(%)	MEMUTIL(%)	HANDLES COUNT	DISK IO READ(KB/SEC)	DISK IO WRITE(KB/SEC)	DISK IOPS(OPERATIONS/SEC)	PAGE FAULTS(FAULTS/SEC)	WEBSITE TITLE	
Oct 25, 2013 18:41:10										
	4404	0	0.7658	544	0	0	0	0	Yahoo India - Windows Internet Explorer	-
	4188	0	0.4282	527	0	0	0	0	-	-
	6132	0	0.465	450	0	0	0	0	Google - Windows Internet Explorer	https:/, gws_rd=

Figure 3.11: The detailed diagnosis of the Running browser instances measure

The detailed diagnosis of the **Recent web sites** measure reveals the names and URLs of the web sites that are being accessed using a browser.

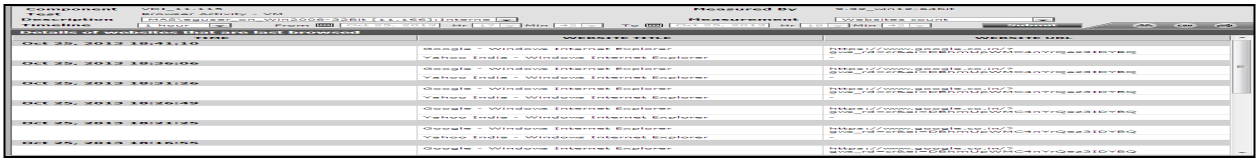

Component VDI_11.115
Test Browser Activity - VM
Description MAS\eguser_on_Win2008-32Bit [11.166]:Interne
Timeline 1 hour From Oct 25, 2013 Hr 17 Min 42 To Oct 25, 2013 Hr 18 Min 42
Measured By 9.32_win12-64bit
Measurement Websites count

Details of websites that are last browsed

TIME	WEBSITE TITLE	WEBSITE URL
Oct 25, 2013 18:41:10		
	Google - Windows Internet Explorer	https://www.google.co.in/?gws_rd=cr&ei=DBhmUpWMC4nYrQez3IDYBQ
	Yahoo India - Windows Internet Explorer	-
Oct 25, 2013 18:36:06		
	Google - Windows Internet Explorer	https://www.google.co.in/?gws_rd=cr&ei=DBhmUpWMC4nYrQez3IDYBQ
	Yahoo India - Windows Internet Explorer	-
Oct 25, 2013 18:31:26		
	Google - Windows Internet Explorer	https://www.google.co.in/?gws_rd=cr&ei=DBhmUpWMC4nYrQez3IDYBQ
	Yahoo India - Windows Internet Explorer	-
Oct 25, 2013 18:26:49		
	Google - Windows Internet Explorer	https://www.google.co.in/?gws_rd=cr&ei=DBhmUpWMC4nYrQez3IDYBQ
	Yahoo India - Windows Internet Explorer	-
Oct 25, 2013 18:21:25		
	Google - Windows Internet Explorer	https://www.google.co.in/?gws_rd=cr&ei=DBhmUpWMC4nYrQez3IDYBQ
	Yahoo India - Windows Internet Explorer	-
Oct 25, 2013 18:16:55		
	Google - Windows Internet Explorer	https://www.google.co.in/?gws_rd=cr&ei=DBhmUpWMC4nYrQez3IDYBQ

Figure 3.12: The detailed diagnosis of the Recent web sites measure

3.2.7 Personal vDisk – VM Test

The personal vDisk retains the single image management of pooled and streamed desktops while allowing people to install applications and change their desktop settings.

Unlike traditional Virtual Desktop Infrastructure (VDI) deployments involving pooled desktops, where users lose their customizations and personal applications when the administrator alters the base virtual machine (VM), deployments using personal vDisks retain those changes. This means administrators can easily and centrally manage their base VMs while providing users with a customized and personalized desktop experience.

Personal vDisks provide this separation by redirecting all changes made on the user's VM to a separate disk (the personal vDisk) attached to the user's VM. The content of the personal vDisk is blended at runtime with the content from the base VM to provide a unified experience. In this way, users can still access applications provisioned by their administrator in the base VM.

But, what happens if a personal vDisk runs out of space? Simple! Users will no longer be able to hold on to their customizations, allowing them access to only the base VM and the applications installed therein! This outcome beats the entire purpose of having personal vDisks! If this is to be avoided, then administrators should continuously monitor the usage of the personal vDisks, proactively detect a potential space crunch, determine what is causing the rapid erosion of space on the personal vDisk, and fix the root-cause, before desktop users complain. This is where the **Personal vDisk – VM** test helps.

For each VM on a XenServer, this test tracks the status and space usage of its personal vDisk and promptly reports errors / abnormal space usage. This way, administrators can accurately identify personal vDisks with very limited space, which VM such personal vDisks are associated with, and what is consuming too much disk space – user profiles? Or user applications?

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every *user_on_VM*

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

    - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
    - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
    - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
    - Then, click the **Update** button to save the changes.

    Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable

the eG agent to communicate with the XenServer via HTTPS:

- Obtain the server-certificate for the XenServer
- Import the server-certificate into the local certificate store of the eG agent

For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by

default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **AGGREGATE USER SESSIONS** -This flag is closely related to the **REPORT BY USER** flag. Since the **REPORT BY USER** flag is set to **No** by default for a Citrix XenServer, this test will, by default, ignore the status of the **AGGREGATE USER SESSIONS** flag while monitoring that server. In case of the Citrix XenServer - VDI on the other hand, the **REPORT BY USER** flag is set to **Yes** by default. Therefore, the status of the **AGGREGATE USER SESSIONS** flag gains significance in the case of this model. By default, the **AGGREGATE USER SESSIONS** flag is set to **No**. This implies that if a single user is currently logged into multiple guests, then this test, by default, will report a set of measures for every *username* on *guestname*. On the other hand, if the status of this flag is changed to **Yes**, then, this test will report a set of (aggregated) measures for every distinct user to the virtual desktop environment. In other words, this test will report measures that are aggregated across all the currently active sessions for a user, spanning multiple VMs.

15. **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 *virtualmachine* 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.

16. **REPORT POWERED ON** - You can set the **REPORT POWERED ON** status to **Yes**, so that the test reports an additional measure, Is VM powered on?, revealing whether a guest OS is currently running or not. The default status of this flag is set to Yes for a Citrix XenServer. For a Citrix XenServer – VDI component on the other hand, this flag is set to No by default. This is because, in such environments, the virtual desktops will be in the powered-off state most of the time.

17. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports *80* or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

18. **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.

Measurements made by the test

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td><b>Personal vDisk service status:</b></td><td>Indicates whether Citrix Personal vDisk service is running or not on this VM.</td><td></td><td>The values that this measure can report and their corresponding numeric values have been discussed in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Stopped</td><td>0</td></tr><tr><td>Running</td><td>1</td></tr><tr><td>Not installed</td><td>2</td></tr></table>
<b>Note:</b>
By default, this test reports the <b>Measures Value</b>s listed in the table above to indicate the status of the Personal vDisk service. In the graph of this measure however, the same will be represented using the numeric equivalents.</td></tr>
<tr><td><b>Recompose status:</b></td><td>Indicates the status of the initially provisioned disk or the updated image.</td><td>Number</td><td>Use the detailed diagnosis of this measure to know for which VM the initial personal vDisk provisioning or image update were unsuccessful and why. The VM can be in one of the following states:<ul><li><b>OK</b> – The initial provisioning or last image update was successful.</li><li><b>Disk Init</b> – This is the first time that the personal vDisk has started or been resized. It is being initialized and</li></ul></td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			partitioned by the service. - **Disk Format** – The personal vDisk is being formatted. - **Updating** – The initial provisioning or an image update is in progress. - **Error (Disk Discovery)** – An error state. An error occurred while discovering the personal vDisk. - **Error (Disk Init)** – An error state. An error occurred while partitioning or formatting the personal vDisk. - **Error (Sys Init)** – An error state. An error occurred while starting the Personal vDisk Service or configuring the personal vDisk. - **Error (Update)** – An error state. An error occurred during the initial provisioning or the last image update. - **Unknown** – An error state. An error occurred but the cause is unknown.
Space used by user applications:	Indicates the amount of space used by applications installed on the personal vDisk attached to this VM.	MB	Personal vDisks have two parts, which use different drive letters and are by default equally sized. One part comprises a Virtual Hard Disk file (a .vhd file). This contains items such as applications installed in C:\Program Files. By default, this part uses drive V: but is hidden from users. These measures indicate how much space has been allocated to this .vhd file and how much of the allocated space has been utilized by user applications contained in this file. A high value for the Space used by user applications and Space utilized by user applications measures is indicative of

Measurement	Description	Measurement Unit	Interpretation
			excessive space used by user applications. You can compare the value of these measures across VMs to know which user to which VM has utilized too much space reserved for user applications on the personal vDisk. If the value of the Space utilized by user applications measure grows close to 100% for any VM, it implies that potentially, the user to that VM will not be able to install any applications on the personal vDisk; nor access any applications.
Space allocated for user applications:	Indicates the amount of space allocation for storing user applications on the personal vDisk attached to this VM.	MB	
Space utilized by user applications:	Indicates the percentage of allocated space used by applications installed on the personal vDisk attached to this VM.	Percent	
Space used by user profiles:	Indicates the amount of space used for storing user profiles on the personal vDisk attached to this VM.	MB	Personal vDisks have two parts, which use different drive letters and are by default equally sized. One part comprises C:\Users (in Windows 7) or C:\Documents and Settings (in Windows XP). This contains user data, documents, and the user profile. By default this uses drive P:. These measures indicate how much space has been allocated to user profiles and how much of the allocated space has been utilized by user profiles. A high value for the Space used by user profiles and Space utilized by user profiles measures is indicative of excessive space used by user profiles. You can compare the value of these measures across VMs to know which VM's user profiles are consuming the maximum space on the personal vDisk. If the value of the Space utilized by user profiles measure grows close to 100% for any VM, it implies that potentially, the user to that VM will not be able to store/access any more documents or

Measurement	Description	Measurement Unit	Interpretation
			user data on the personal vDisk .
Space allocated for user profiles:	Indicates the amount of space allocated for storing user profiles on the personal vDisk attached to this VM.	MB	
Space utilized by user profiles:	Indicates the percentage of allocated space that has been used up by user profiles on the personal vDisk attached to this VM.	Percent	
Free space:	Indicates the amount of unused space on the personal vDisk attached to this VM.	MB	Ideally, the value of this measure should be high. You can compare the value of this measure across VMs to know which VM's personal vDisk has the least free space. You may then want to resize that personal vDisk to accommodate more data.
Total size:	Indicates the total size of the personal vDisk attached to this VM.	MB	The minimum size of a Personal vDisk is 3 GB, however a size of 10 GB is recommended.
Space utilized:	Indicates the percentage of space in the personal vDisk attached to this VM that is currently used.	Percent	A consistent increase in the value of this measure is a cause for concern, as it indicates a gradual erosion of free space in the personal vDisk of a VM. By comparing the value of this measure across VMs, you can identify which VM's personal vDisk is running out of space! Once the VM with the space-hungry vDisk is isolated, you may want to compare the value of the **Space utilized by user applications** and **Space utilized by user profiles** measures of that VM, to clearly understand what is occupying too much

Measurement	Description	Measurement Unit	Interpretation
			space in the personal vDisk – is it the user profiles? Or is it the user applications? Based on this inference, you can figure out which drive partition of the personal vDisk has limited free space, and can decide between freeing up space in that partition or allocating more space to the personal vDisk itself.

3.2.8 Virtual Desktop Session Start-up Details Test

Figure 3.13 depicts a typical user logon process to a virtual desktop via XenDesktop broker.

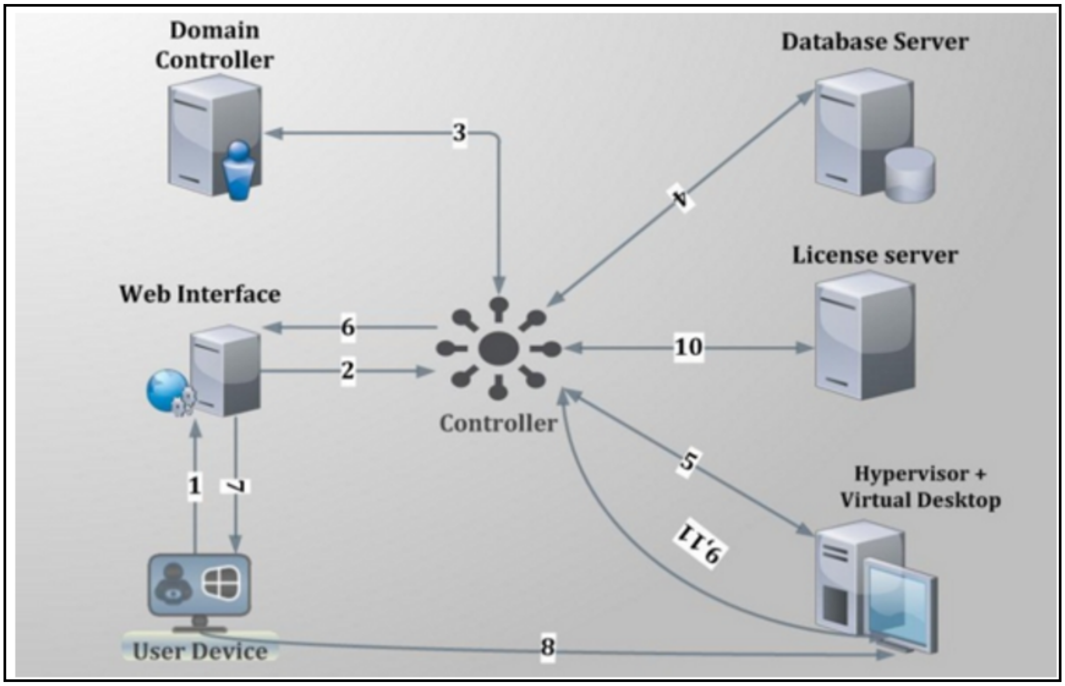

Figure 3.13: Citrix user logon process

The process depicted by Figure 3.13 above has been described below:

1. User provides his/her credentials to the web interface.
2. Web interface forwards the credentials to controller for verification process.

3. Delivery controller transfers these credentials to the domain controller to check if the user is present in the active directory.
4. Once it gets the successful confirmation from AD then controller communicates with site database to check what type of virtual desktop is available for current user.
5. Controller then interacts with the hypervisor layer to gather information about the availability of virtual desktop.
6. Controller then passes the ICA file for user and all the connection information is present inside ICA file so that client can establish the connection.
7. After all the process is complete, the user is assigned the virtual desktop.
8. The user then establishes a connection with the assigned virtual desktop.
9. The virtual desktop again communicates with controller for verification of licensing.
10. Controller checks for license from license server about what type of license is available for user in this current session. License server then communicates back with controller providing the licensing information.
11. Information obtained from license server is then passed to the virtual desktop.

From the discussion above, it can be inferred that login processing happens at two different places – at the delivery controller, and inside the virtual desktop. While login, authentication, and application brokering happen on the delivery controller, session creation and setup happens inside the virtual desktop. A problem in any of these places can result in a poor user experience. Inevitably, these issues result in service desk calls and complaints that “Citrix is slow.” Diagnosing login problems has traditionally been a difficult, time-consuming, manual process due to the large number of steps involved. The key to resolving user experience issues therefore, lies in tracking each user’s sessions end-to-end, ascertaining the time spent by the session at each step of the logon process – be it on the delivery controller or on the virtual desktop– and accurately identifying where and at what step of the logon process, the slowdown occurred.

To determine the time taken by the entire logon process of a user, isolate logon slowness, and understand where the process was bottlenecked – whether on the delivery controller or on the XenApp server – use the **User Logon Performance** test mapped to the **Citrix XA/XD Site** component. If the **User Logon Performance** test reveals a problem in session start-up on the virtual desktop, then use the **Virtual Desktop Session Start-up Details** test.

With the **Virtual Desktop Session Start-up Details** test, administrators can receive deep visibility into the **virtual desktop end** of the Citrix logon process. This test takes an administrator into the virtual desktop, reveals the users who are currently logged on to the virtual desktop, and accurately reports the average time it took for the sessions of each user to start inside the virtual desktop. This way, administrators can rapidly identify which user’s sessions are experiencing undue start-up delays.

In addition, the test also provides a break-up of the session start-up duration. This way, the test precisely pinpoints where the delay occurred – when user credentials were obtained? when credentials were validated? during profile loading? during login script execution? when mapping drives or creating printers?

For this purpose, the test categorizes its metrics into *client start-up metrics* and *server start-up metrics*.

The *client start-up metrics* are concerned with timing the operations that occur from the point when the user requests for access to a virtual desktop to the point at which a connection to the virtual desktop is established. While onnection-brokering mechanisms involve components that are not on the physical client device, the tasks these systems perform have a direct impact on the performance of the connection start-up and are recorded as part of the client-side process.

The *server start-up metrics* are concerned with timing the operations that occur when creating a new session on the virtual desktop. This includes user authentication, client device mapping, profile loading, login scripts execution, and finally, starting the user's desktop.

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every *user_on_VM*

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

    - Modify the target XenServer’s configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
    - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
    - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right

to connect to the XenServer console via SSH.

- Then, click the **Update** button to save the changes.

Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.
6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.
7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

- **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

14. **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 *virtualmachine* 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.

15. **REPORT POWERED ON** - You can set the **REPORT POWERED ON** status to **Yes**, so that the test reports an additional measure, Is VM powered on?, revealing whether a guest OS is currently running or not. The default status of this flag is set to Yes for a Citrix XenServer. For a Citrix XenServer – VDI component on the other hand, this flag is set to No by default. This is because, in such environments, the virtual desktops will be in the powered-off state most of the time.

16. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports *80* or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

17. **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.

Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
User sessions:	Indicates the number of sessions currently open for this user on this virtual desktop.	Number	Use the detailed diagnosis of this measure to view the complete details of this user's session. Such details includes the name and IP address of the client from which the session was launched, when session creation started, and when it ended. With the help of this information, administrators can quickly understand if the session took too long to get created.
Session start- up duration:	Indicates the time taken by this user to complete session start-up inside this virtual desktop.	Secs	Compare the value of this measure across users to know which user's sessions took the longest to start on the virtual desktop. To know what is causing this 'slowness', compare the values reported by all the other 'duration' measures of this test for that user on that virtual desktop. This will quickly lead you to where that user's session start-up is spending the maximum time.
Profile load duration:	Indicates the time taken to load this user's profile.	Secs	If the user's Session start-up duration is high, you may want to compare the value of this measure with that of the other 'duration' measures reported for this user to figure out if a delay in profile

Measurements made by the test	Measurement	Measurement Unit	Interpretation
			loading is what is really ailing that user's logon experience with this virtual desktop. One of the common reasons for high profile load time is the large size of the user profile.
Group policy processing duration:	Indicates the time taken by this user's session to process group policies.	Secs	If a user's Session start-up duration is high, you may want to compare the value of this measure with that of the other 'duration' measures reported for this user to figure out if a delay in group policy processing is what is really ailing that user's logon experience with this virtual desktop. In such a case, you can also use the detailed diagnosis of this measure to figure out the names of the group policy client- side extensions (CSE), the time each CSE took to run, the status of every CSE, and errors (if any) encountered by each CSE. Using these in- depth metrics, Citrix administrators can accurately pinpoint which CSE is impeding speedy group policy processing. **Note:** Detailed diagnostics will be available for this measure only if the eG VM Agent is deployed on the virtual desktops and the inside view using parameter of this test is set to eG VM Agent. Logon performance improves when fewer Group Policies are applied. Merge GPOs when possible instead of having multiple GPOs.
Login script execution duration:	Indicates the time taken for the login script to execute	Secs	If a user complains of slowness, then, you can compare the value of this

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	for this user.		measure with that of the other 'duration' measures of that user to figure out what could have really caused the slowness.
Start- up client duration:	This is the high-level client-side connection start- up metric. It starts at the time of the request (mouse click) and ends when the connection between this user's client device and the virtual desktop has been established.	Secs	When any user complains of slowness when trying to logon to a virtual desktop, you may want to compare the value of this measure with that of the Session start-up server duration measure of that user to know whether a client-side issue or a server-side issue is responsible for the slowness he/she is experiencing with that virtual desktop. If this comparison reveals that the Start-up client duration of the user is high, it indicates a client- side issue that is causing long start times. In this case therefore, compare the value of the client start- up metrics such as the Application enumeration client duration, Configuration file download client duration, Credentials obtention client duration, ICA file download client duration, Launch page web server duration, Name resolution client duration, Name resolution web server duration, Session look- up client duration, Session creation client duration, and Ticket response web server duration to know what client-side issue is causing the Start- up client duration to be high.
Back-up URL client count:	**This measure is relevant when the Citrix Receiver is the session launch mechanism.** It records the number of back- up URL retries before a successful	Number	If this metric has a value higher than 1, it indicates that the Web Interface server is unavailable and the Citrix Receiver is attempting to connect to back-up Web Interface servers to launch the virtual desktop.

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	launch. Note that this is the only start-up metric that is a measure of attempts, rather than time duration.		A value of 2 means that the main Web Interface server was unavailable, but the Citrix Receiver managed to launch the virtual desktop successfully using the first back-up server that it tried. A value higher than 2 means that multiple Web Interface servers are unavailable. Probable reasons for the non-availability of the Web Interface servers include (in order of likelihood): • Network issues between the client and the server. So the administrator should make sure that the Web Interface server is on the network and accessible to the clients. • An overloaded Web Interface server that is not responding (or has crashed for another reason). Try to log on to the server and check the Windows Performance Monitor/Task Manager to see how much memory is in use and so on. Also, review the Event Logs to see if Windows logged any serious errors.
Application enumeration client duration:	**This measure is relevant when the Citrix Receiver is the session launch mechanism.** It measures the time needed by this user's session to retrieve the list of applications from the Web Interface service.	Secs	If the Start-up client duration measure reports a high value for a user, then compare the value of this measure with that of the other client-side metrics such as Configuration file download client duration, Credentials obtention client duration, ICA file download client duration, Launch page web server duration, Name resolution client duration, Name resolution web server duration, Session look- up client

Measurements made by the test	Measurement	Measurement Unit	Interpretation
			duration, Session creation client duration, and Ticket response web server duration to know whether/not slowness in application enumeration is the precise reason why it took the user a long time to establish a session with the virtual desktop.
Configuration file download client duration:	**This measure is relevant when the Citrix Receiver is the session launch mechanism.** It measures the time this user's session took to retrieve the configuration file from the XML broker.	Secs	If the Start-up client duration measure reports a high value for a user, then compare the value of this measure with that of the other client-side metrics such as Application enumeration client duration, Credentials obtention client duration, ICA file download client duration, Launch page web server duration, Name resolution client duration, Name resolution web server duration, Session look- up client duration, Session creation client duration, and Ticket response web server duration to know whether/not slowness in retrieving the configuration file from the XML server is the precise reason why it took the user a long time an ICA session with the XenApp server.
Credentials obtention client duration:	**This measure is relevant when the Citrix Receiver is the session launch mechanism.** It measures the time required by this user's session to obtain the user credentials.		Note that COCD is only measured when the credentials are entered manually by the user. Because this metric may be artificially inflated if a user fails to provide credentials in a timely manner, it is subtracted from the Start-up client duration. However, in the event that the user manually inputs the credentials, and the value of this measure is higher than that of all the other client start-up metrics that this test reports, it is a clear

Measurements made by the test	Measurement	Measurement Unit	Interpretation
			indicator that any connection delay that the user may have experienced is owing to slowness in obtaining user credentials.
ICA file download duration:	**This measure is relevant when the Citrix Receiver is the session launch mechanism.** This is the time it takes for this user's client to download the ICA file from the web server.	Secs	The overall process here is: 1. The user clicks on application icon. The user's browser requests the Web Interface launch page. The Web Interface launch page receives the request and starts to process the launch, communicating with the virtual desktop and potentially other components such as Secure Ticket Authority (STA). The Web Interface generates ICA file data. The Web Interface sends the ICA file data back to the user's browser. The browser passes ICA file data to the client. This measure represents the time it takes for the complete process (step 1 to 6). The measure stops counting time when the client receives the ICA file data. The Launch page web server duration measure on the other hand, covers the Web server portion of the process (that is, steps 3 and 4). If the ICA file download duration is high, but the Launch page web server duration is normal, it implies that the server-side processing of the launch was successful, but there were communication issues between the

Measurements made by the test	Measurement	Measurement Unit	Interpretation
			client device and the Web server. Often, this results from network trouble between the two machines, so investigate potential network issues first.
Launch page web server duration:	**This measure is relevant when the Web Interface is the session launch mechanism.** It measures the time needed by this user's session to process the launch page (launch. aspx) on the Web Interface server.	Secs	If the value of this measure is high, it indicates at a bottleneck on the Web Interface server. Possible causes include: • High load on the Web Interface server. Try to identify the cause of the slow down by checking the Internet Information Services (IIS) logs and monitoring tools, Task Manager, Performance Monitor and so on. • Web Interface is having issues communicating with the other components. Check to see if the network connection between Web Interface and virtual desktop is slow. If the Web server seems okay, consider reviewing the virtual desktop for problems.
Name resolution client duration:	This is the time it takes the XML service to resolve the name of a published application to an IP address.	Secs	This metric is collected when a client device directly queries the XML Broker to retrieve published application information stored in IMA. This measure is only gathered for new sessions since session sharing occurs during startup if a session already exists. When this metric is high, it indicates the XML Broker is taking a lot of time to resolve the name of a published application to an IP address. Possible causes include a problem on the client,

Measurements made by the test	Measurement	Measurement Unit	Interpretation
			issues with the XML Broker, such as the XML Broker being overloaded, a problem with the network link between the two, or a problem in IMA. Begin by evaluating traffic on the network and the XML Broker.
Name resolution web server duration:	**This measure is relevant when the Citrix Receiver is the session launch mechanism.** It is the time it takes the XML service to resolve the name of this virtual desktop to its IP address.	Secs	When this metric is high, there could be an issue with the Web Interface server or the Citrix Receiver, the XML Service, the network link between the two, or a problem in IMA. Like the Name resolution client duration measure, this metric indicates how long it takes the XML service to resolve the name of a virtual desktop to its IP address. However, this metric is collected when a Web Interface site is performing this process on behalf of a launch request it has received from either the Citrix Receiver or from a user clicking a Web Interface page icon.
Session look- up client duration:	Indicates the time this user's session takes to query every ICA session to host the requested published application.	MSecs	The check is performed on the client to determine whether the application launch request can be handled by an existing session. A different method is used depending on whether the session is new or shared.
Session creation client duration:	Indicates the new session creation time.	Secs	In the event of slowness, if the Start-up client duration of a user session is found to be higher than the Session start-up server duration, you may want to compare the value of this measure with all other client start- up measures to determine whether/not session creation is the process that is slowing down the application launch.

Measurements made by the test	Measurement	Measurement Unit	Interpretation
Ticket response web server duration:	**This measure is relevant when the Citrix Receiver or Web Interface is the desktop launch mechanism.** This is the time this user's sessions take to get a ticket (if required) from the STA server or XML service.	Secs	When this metric is high, it can indicate that the Secure Ticket Authority (STA) server or the XML Broker are overloaded.
Reconnect enumeration client duration:	**This measure is relevant when the Citrix Receiver or Web Interface is the desktop launch mechanism.** This is the time it takes this user's client to get a list of reconnections.	Secs	Compare the value of this measure with that of other client start-up metrics for a user to know what is the actual cause for the client start-up delay.
Reconnect enumeration web server duration:	**This measure is relevant when the Citrix Receiver or Web Interface is the desktop launch mechanism.** This is the time it takes the Web Interface to get the list of reconnections for this user from the XML service.	Secs	Compare the value of this measure with that of other client start-up metrics for a user to know what is the actual cause for the client start-up delay.
Session start- up server duration:	This is the high- level server- side connection start-up metric. It includes the time spent on this virtual desktop to perform the entire start- up operation.	Secs	When this metric is high, it indicates that there is a server- side issue increasing session start times. To zero-in on this issue, compare the values of the server start- up metrics such as Session creation server duration, Credentials obtention server duration, Program neighbourhood credentials obtention server duration, Credentials obtention network server duration, Credentials authentication server

Measurements made by the test	Measurement	Measurement Unit	Interpretation
			duration, Profile load server duration, Login script execution server duration, Drive mapping server duration, Drive mapping server duration, and Printer creation server duration
Session creation server duration:	Indicates the time spent by this virtual desktop in creating the session for this user.	Secs	This duration starts when the ICA client connection has been opened and ends when authentication begins. This should not be confused with 'Session start-up server duration'.
Credentials obtention server duration:	Indicates the time taken by this virtual desktop to obtain the credentials of this user.	Secs	This time is only likely to be a significant if manual login is being used and the server- side credentials dialog is displayed (or if a legal notice is displayed before login commences). Because this metric may be artificially inflated if a user fails to provide credentials in a timely manner, it is not included in the Session start-up server duration. However, in the event that the user manually inputs the credentials, and the value of this measure is higher than that of all the other client start-up metrics that this test reports, it is a clear indicator that any connection delay that the user may have experienced is owing to slowness in obtaining user credentials.
Credentials obtentions network server duration:	Indicates the time spent by this virtual desktop performing network operations to obtain credentials for this user.	Secs	This only applies to a Security Support Provider Interface login (a form of pass-through authentication where the client device is a member of the same domain as the server and Kerberos tickets are passed in place of manually entered credentials).

Measurements made by the test	Measurement	Measurement Unit	Interpretation
Program neighbourhood credentials obtention server duration:	Indicates the time needed for this virtual desktop to cause the Program Neighborhood instance running on the client ("Program Neighborhood Classic") to obtain this user's credentials.	Secs	As in the case of the Credentials obtention server duration metric, because this metric may be artificially inflated if a user fails to provide credentials in a timely manner, it is not included in the Session start-up server duration.
Credentials authentication server duration:	Indicates the time spent by this virtual desktop when authenticating the user's credentials against the authentication provider, which may be Kerberos, Active Directory® or a Security Support Provider Interface (SSPI).	Secs	Where server-side issues are causing user experience to deteriorate, you can compare the value of this measure with that of all the other server start- up metrics that this test reports – i.e., Session creation server duration, Credentials obtention server duration, Program neighbourhood credentials obtention server duration, Credentials obtention network server duration, Profile load server duration, Login script execution server duration, Drive mapping server duration, Drive mapping server duration, and Printer creation server duration – to know what is the root-cause of delays in server start-up.
Profile load server duration:	Indicates the time required by this virtual desktop to load this user's profile.	Secs	If this metric is high, consider your Terminal Services profile configuration. Citrix Consulting has found that when customers have logon times greater than 20 seconds, in most cases, this can be attributed to poor profile and policy design. Roaming profile size and location contribute to slow session starts. When a user logs onto a session where Terminal Services roaming profiles and home folders are enabled, the roaming profile contents and access to that folder are mapped during logon, which takes additional resources. In

Measurements made by the test	Measurement	Measurement Unit	Interpretation
			some cases, this can consume significant amounts of the CPU usage. Consider using the Terminal Services home folders with redirected personal folders to mitigate this problem. In general, consider using Citrix Profile management to manage user profiles in Citrix environments. This tool also provides logging capabilities to help isolate profile issues. If you are using Citrix profile management and have slow logon times, check to see if your antivirus software is blocking the Citrix profile management tool.
Login script execution server duration:	Indicates the time needed by this virtual desktop to run this user's login script (s).	Secs	If the value of this measure is abnormally high for any user, consider if you can streamline this user or group's login scripts. Also, consider if you can optimize any application compatibility scripts or use environment variables instead.
Drive mapping server duration:	Indicates the time needed for this virtual desktop to map this user's client drives, devices and ports.	Secs	Make sure that, when possible, your base policies include settings to disable unused virtual channels, such as audio or COM port mapping, to optimize the ICA protocol and improve overall session performance.
Printer creation server duration:	Indicates the time required for this virtual desktop to synchronously map this user's client printers.	Secs	If the configuration is set such that printer creation is performed asynchronously, no value is recorded for this measure as it is does not impact completion of the session start-up. On the other hand, if excessive time is spent mapping printers, it is often the

Measurements made by the test	Measurement	Measurement Unit	Interpretation
			result of the printer autocreation policy settings. The number of printers added locally on the users' client devices and your printing configuration can directly affect your session start times. When a session starts, the virtual desktop has to create every locally mapped printer on the client device. Consider reconfiguring your printing policies to reduce the number of printers that get created - especially if users have a lot of local printers.

3.2.9 Virtual Desktop Sessions Details Test

A user logged into a virtual desktop does not imply active usage of that desktop. In a VDI infrastructure, it is common for users to just log into desktops, and leave them unused for long time periods. Such desktops are a huge resource drain, as they continue to consume resources, regardless of the level of activity on them. Idle users themselves are unproductive resources. Besides, since these users unnecessarily hold on to desktops, users with genuine needs may not have any desktops to work with. If administrators can quickly identify these idle users and the desktops they are logged into, they can rapidly pull the desktops from such users and assign them to users who can use them effectively. The **Virtual Desktop Sessions Details** test turns the spotlight on these idle users. For each user session on a virtual desktop, this test reports the total duration of the session and the percentage of time for which the session was active. The test also reports the total idle time during the session. From these statistics, administrators can accurately identify those users who are wasting the desktops assigned and resources allocated to them.

-

Target of the test : A Citrix XenServer

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every user who is currently logged into a virtual desktop

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.

3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.

4. **XEN USER** - To enable the eG agent to connect to the XenServer API for collecting statistics of interest, this test should login to the XenServer as a **root** user. Provide the name of the **root** user in the XEN USER text box. **Root** user privileges are **mandatory** when monitoring a **XenServer 5.5 (or below)**. However, if you are monitoring **XenServer 5.6 (or above)** and you prefer not to expose the credentials of the **root** user, then, you have the option of configuring a user with **pool-admin** privileges as the **XEN USER**. If you do not want to expose the credentials of a **root/pool-admin user**, then you can configure the tests with the credentials of a **XEN USER** with **Read-only privileges** to the XenServer. However, if this is done, then the **Xen Uptime** test will not run, and the **Xen CPU** and **Xen Memory** tests will not be able to report metrics for the **control domain** descriptor. To avoid such an outcome, do the following before attempting to configure the eG tests with a **XEN USER** who has **Read-only** privileges to the XenServer:

   - Modify the target XenServer's configuration in the eG Enterprise system. For this, follow the Infrastructure -> Components -> Add/Modify menu sequence, pick *Citrix XenServer* as the **Component type**, and click the **Modify** button corresponding to the target XenServer.
   - In the **MODIFY COMPONENT DETAILS** page that then appears, make sure that the **OS** is set to **Xen** and the **Mode** is set to **SSH**.
   - Then, in the same page, proceed to provide the **User** and **Password** of a user who has the right to connect to the XenServer console via SSH.
   - Then, click the **Update** button to save the changes.

   Once this is done, you can configure the eG tests with the credentials of a **XEN USER** with *Read-only* privileges.

5. **XEN PASSWORD** - The password of the specified **XEN USER** needs to be mentioned here.

6. **CONFIRM PASSWORD** - Confirm the **XEN PASSWORD** by retyping it here.

7. **SSL** - By default, the Xen Server is not SSL-enabled. This indicates that by default, the eG agent communicates with the XenServer using HTTP. Accordingly, the **SSL** flag is set to **No** by default. If you configure the XenServer to use SSL, then make sure that the **SSL** flag is set to **Yes**, so that the eG agent communicates with the XenServer using HTTPS. Note that a default SSL certificate comes bundled with every XenServer installation. If you want the eG agent to use this default certificate for communicating with an SSL-enabled XenServer, then no additional configuration is required. However, if you do not want to use the default certificate, then you can generate a self-signed certificate for use by the XenServer. In such a case, you need to explicitly follow the broad steps given below to enable the eG agent to communicate with the XenServer via HTTPS:

   - Obtain the server-certificate for the XenServer
   - Import the server-certificate into the local certificate store of the eG agent

   For a detailed discussion on each of these steps, refer to the *Troubleshooting* section of this document.

8. **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*.

9. **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 **2.7.2**.

   - **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 PASSWORD**s 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 **2.5.1.1**.

- **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.

10. **IGNORE VMS INSIDE VIEW** - Administrators of some high security XenServer 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 a XenServer 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.

11. **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.

12. **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.

13. **WEBPORT** - By default, in most virtualized environments, the XenServer listens on port 80 (if not SSL-enabled) or on port 443 (if SSL-enabled). This implies that while monitoring an SSL-enabled XenServer, the eG agent, by default, connects to port 443 of the server to pull out metrics, and while monitoring a non-SSL-enabled XenServer, the eG agent connects to port 80. Accordingly, the **WEBPORT** parameter is set to *80* or *443* depending upon the status of the **SSL** flag. In some environments however, the default ports 80 or 443 might not apply. In such a case, against the **WEBPORT** parameter, you can specify the exact port at which the XenServer in your environment listens so that the eG agent communicates with that port.

14. **REPORT BY USER** - While monitoring a *Citrix XenServer*, the **REPORT BY USER** flag is set to **No** by default, indicating that by default, the guest operating systems on the XenServer are identified using

the hostname specified in the operating system. On the other hand, while monitoring a *Citrix XenServer - VDI*, this flag is set to **Yes** by default; this implies that in case of the XenServer 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*.

15. **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 *virtualmachine* 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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total time in session:	Indicates the time that has elapsed since this user logged into this desktop.	Mins	
Active time in last measure period:	Indicates the percentage of time in the last measurement period during which this user actively used this desktop.	Percent	Ideally, the value of this measure should be 100%. A low value for this measure denotes a high level of inactivity recently.
Time since last activity:	Indicates the time that has elapsed since this user performed an action on this desktop.	Mins	A high value for this measure indicates that the user has been idle for a long time. Compare the value of this measure across users to know which user has been idle for the longest time.
Total idle time in session:	Indicates the total time for which this user was idle during the session.	Mins	If the value of this measure is the same as the value of the *Total time in session* measure for a user, it means that the user has been idle throughout the session. If the value of this measure is close to the value of the *Total time in session* measure for a user, it implies that the user has been idle for a long time.

Measurement	Description	Measurement Unit	Interpretation
			If the value of this measure is much lesser than the value of the *Total time in session* measure for a user, it means that the user has been active for most part of the session.

Conclusion

This document has clearly explained how eG Enterprise monitors XenServers. We can thus conclude that eG Enterprise, with its ability to provide in-depth insight into the performance of virtualized Xen infrastructures, is the ideal solution for monitoring such environments. For more information on eG Enterprise, please visit our web site at www.eginnovations.com or write to us at sales@eginnovations.com.