



## Monitoring Cloud Desktops

eG Innovations Product Documentation

[www.eginnovations.com](http://www.eginnovations.com)

**eG**  
Total Performance Visibility

# Table of Contents

---

CHAPTER 1: INTRODUCTION .....	1
CHAPTER 2: HOW DOES EG ENTERPRISE MONITOR CLOUD DESKTOPS? .....	2
2.1 Pre-Requisites for monitoring Cloud Desktops .....	2
2.2 Licensing .....	4
CHAPTER 3: HOW TO MONITOR CLOUD DESKTOPS USING EG ENTERPRISE? .....	5
3.1 Installing the eG VM Agent .....	5
3.1.1 Using the Executable Provided by eG Enterprise .....	6
3.1.2 Using SCCM .....	9
3.1.3 Silent Installation of the eG VM Agent .....	25
3.2 Configuring the eG VM Agent to Communicate with the eG Manager .....	26
3.3 Enabling a TCP Port on the eG Remote Agent .....	28
3.4 Managing a Cloud Desktops Component in eG Enterprise .....	29
CHAPTER 4: MONITORING THE CLOUD DESKTOPS COMPONENT .....	31
4.1 The Outside View of VMs Layer .....	32
4.1.1 VM Logins - Cloud Test .....	32
4.1.2 Virtual Machines - Cloud Test .....	34
4.2 The Inside View of Desktops Layer .....	37
4.2.1 Disk Activity - VM Test .....	38
4.2.2 Disk Space - VM Test .....	46
4.2.3 GPU - VM Test .....	47
4.2.4 System Details - VM Test .....	70
4.2.5 Memory Usage - VM Test .....	74
4.2.6 Handles Usage - VM Test .....	78
4.2.7 PCoIP Session - VM Test .....	81
4.2.8 Blast Session - VM Test .....	84
4.2.9 RemoteFX User Experience - VM Test .....	86
4.2.10 Tcp Traffic - VM Test .....	91
4.2.11 Uptime - VM Test .....	93
4.2.12 Tcp - VM Test .....	95
4.2.13 Virtual Desktop Sessions Details Test .....	97
4.2.14 Windows Network Traffic - VM Test .....	99
4.2.15 Windows Memory - VM Test .....	101
4.2.16 Windows Security Center Status - VM Test .....	105
4.2.17 Windows Services - VM Test .....	108
4.2.18 Desktop's HDX Channel Test .....	110
4.2.19 Personal vDisk – VM Test .....	117

---

4.2.20 Disk Alignment - VM Test .....	124
4.2.21 Page File - VM Test .....	128
4.2.22 Windows Update Details - VM Test .....	130

# Table of Figures

---

Figure 1.1: How DaaS works .....	1
Figure 2.1: The eG Remote Agent - eG VM Agent communication when monitoring desktops on a hypervisor .....	3
Figure 2.2: The eG VM Agent - eG Remote Agent communication when monitoring desktops on the cloud .....	4
Figure 3.1: Welcome screen of the eG VM Agent installation wizard .....	6
Figure 3.2: Accepting the license agreement .....	7
Figure 3.3: Specifying the install directory of the eG VM Agent .....	7
Figure 3.4: Specifying the VM agent port .....	8
Figure 3.5: A summary of your specifications .....	8
Figure 3.6: Finishing the installation .....	9
Figure 3.7: Specifying package information .....	12
Figure 3.8: Specifying the command to run .....	13
Figure 3.9: Selecting the platform on which the package should be deployed .....	14
Figure 3.10: A summary of the package specifications .....	15
Figure 3.11: The message that appears upon successful package creation .....	16
Figure 3.12: Selecting the eGVMAgent package for deployment .....	17
Figure 3.13: Deploying the chosen package .....	17
Figure 3.14: Selecting the collection on which the package is to be deployed .....	18
Figure 3.15: Selecting the distribution points .....	19
Figure 3.16: Specifying an action and purpose .....	20
Figure 3.17: Scheduling the package deployment .....	20
Figure 3.18: Specifying the user experience .....	22
Figure 3.19: Specifying how to run the content for the program .....	23
Figure 3.20: Viewing a summary of the settings .....	23
Figure 3.21: Completion of deployment wizard .....	24
Figure 3.22: Viewing the status of the deployment .....	24
Figure 3.23: Adding a Cloud Desktops component .....	29
Figure 4.1: Layer model of the Cloud Desktop component .....	31
Figure 4.2: The tests mapped to the Outside View of VMs layer .....	32
Figure 4.3: The detailed diagnosis of the New logins measure .....	34
Figure 4.4: The detailed diagnosis of the Total VMs reported in the last hour measure .....	37
Figure 4.5: The detailed diagnosis of the VMs with users measure .....	37
Figure 4.6: The tests mapped to the Inside View of Desktops layer .....	38
Figure 4.7: The detailed diagnosis of the Percent virtual busy measure .....	42
Figure 4.8: Configuring a VM test .....	43
Figure 4.9: The VM user configuration page .....	43
Figure 4.10: Adding another user .....	44
Figure 4.11: Associating a single domain with different admin users .....	45

---

Figure 4.12: The test configuration page displaying multiple domain names, user names, and passwords .....	46
Figure 4.13: Dedicated GPU Technology .....	49
Figure 4.14: Shared vGPU Technology .....	50
Figure 4.15: Custom check box .....	66
Figure 4.16: Selecting the NVIDIA WMI check box .....	67
Figure 4.17: NVIDIA Enterprise Management Toolkit (NVWMI) .....	68
Figure 4.18: NVIDIA WMI minimum driver requirements .....	69
Figure 4.19: Checking the Control Panel for NVIDIA WMI .....	70
Figure 4.20: The top 10 CPU consuming processes .....	73
Figure 4.21: The detailed diagnosis of the Handles used by processes measure .....	80
Figure 4.22: The detailed diagnosis of the Processes using handles above limit in VM measure .....	81
Figure 4.23: The SAN, VMFS, and NTFS blocks .....	124
Figure 4.24: Unaligned partitions .....	124
Figure 4.25: Aligned partitions .....	125
Figure 4.26: The detailed diagnosis of the Disk partition alignment status measure .....	127

# Chapter 1: Introduction

A Cloud Desktop is a “web based desktop” that you can access from any suitable device connected to the Internet. Such desktops are the mainstay of Desktop-as-a-Service (DaaS) offerings. In a typical DaaS offering, desktop operating systems run inside virtual machines on servers in a cloud provider's (eg., AWS Workspaces, Microsoft Azure, etc.) data center; these desktop operating systems are called Cloud Desktops. All the necessary support infrastructure for DaaS, including storage and network resources, also lives in the cloud. The DaaS provider then streams the cloud desktops over a network to a customer's endpoint devices, where end users may access them through client software or a web browser.

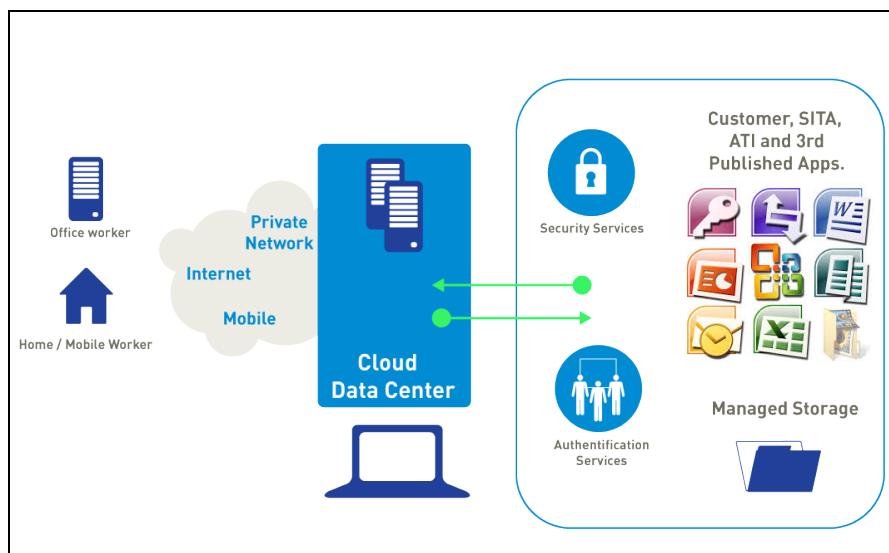


Figure 1.1: How DaaS works

Cloud desktops encounter similar performance problems as do on-premises VDI. Slow logon, poor user experience complaints, unsatisfactory session performance, and resource utilization bottlenecks often hamper cloud desktop delivery. To maximize user satisfaction and earn quick and significant returns on their cloud investments, administrators should monitor the resource usage of and user activity on cloud desktops, proactively detect abnormalities, and promptly fix them. This is where eG Enterprise helps!

eG Enterprise provides a specialized Cloud Desktop model which provides a 360 degree view of the performance of desktops on a cloud. This document will discuss how eG Enterprise monitors VDI on the cloud and what performance insights it provides:

## Chapter 2: How Does eG Enterprise Monitor Cloud Desktops?

eG Enterprise monitors cloud desktops in an agentless manner.

### 2.1 Pre-Requisites for monitoring Cloud Desktops

- For internal monitoring of the individual Windows desktops on the cloud, a light-weight eG VM Agent has to be installed on each of the Windows desktops.

Typically, when **monitoring the Windows desktops on hypervisors** such as VMware vSphere, Citrix XenServer, etc., the eG VM Agent installed on each of the Windows desktops periodically collects and caches the metrics on the desktops. Then, an eG remote agent installed on a remote Windows/Linux host in the environment communicates with the hypervisor to automatically discover the IP address and host name of every virtual desktop on that hypervisor. With this information, the remote agent then polls the eG VM Agent on each Windows desktop to collect the cached metrics. The metrics so collected are then sent to the eG manager.

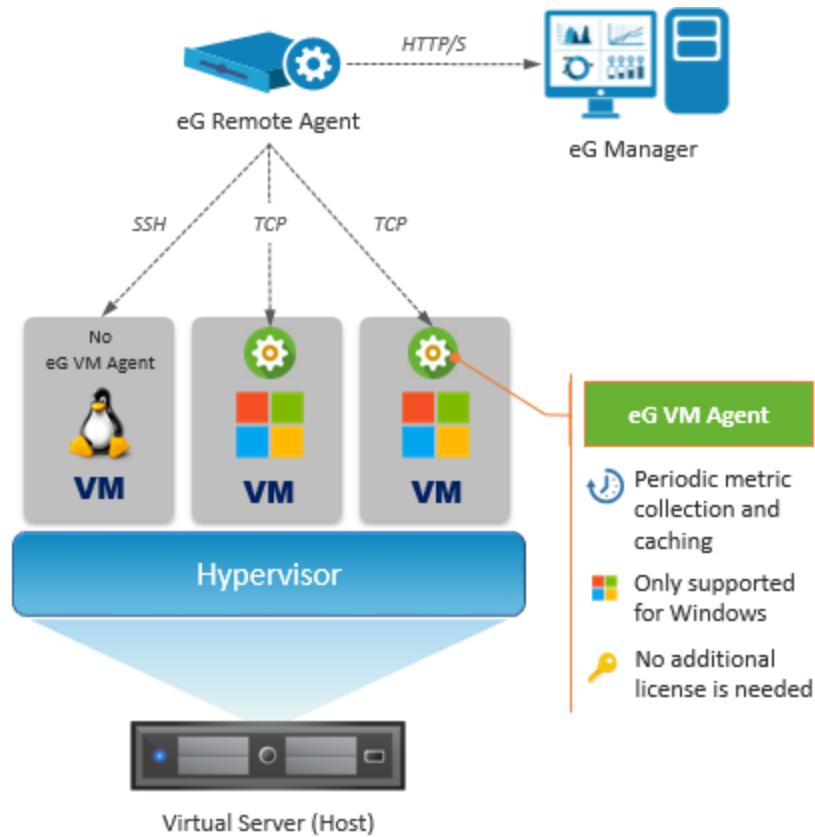


Figure 2.1: The eG Remote Agent - eG VM Agent communication when monitoring desktops on a hypervisor

In the case of a Cloud desktop too, an eG remote agent on a remote Windows/Linux host is used for monitoring. However, unlike the hypervisor scenario, here, the remote agent does not initiate any communication with the eG VM Agent on the individual Windows desktops on the cloud. This is because, this remote agent has no knowledge of the hypervisor on which the cloud desktops are operating. Therefore, it can neither discover the IP address/host name of the cloud desktops, nor can it communicate with them to collect the 'inside view' metrics.

This is why, the eG VM Agent on a Windows cloud desktop has been specifically engineered to initiate all communication. The eG VM Agent on a cloud desktop first communicates with the eG manager to know which remote agent has been assigned to that Windows desktop. Then, the eG VM Agent pushes the cached 'inside view' metrics to that remote agent via TCP. To enable this communication, a TCP port has to be opened on the eG remote agent. If the remote agent is behind a firewall, then make sure that the firewall is configured to allow one-way communication from the eG VM agent to the remote agent. Upon receipt of the 'inside view' metrics, the remote agent sends the metrics to the eG manager.

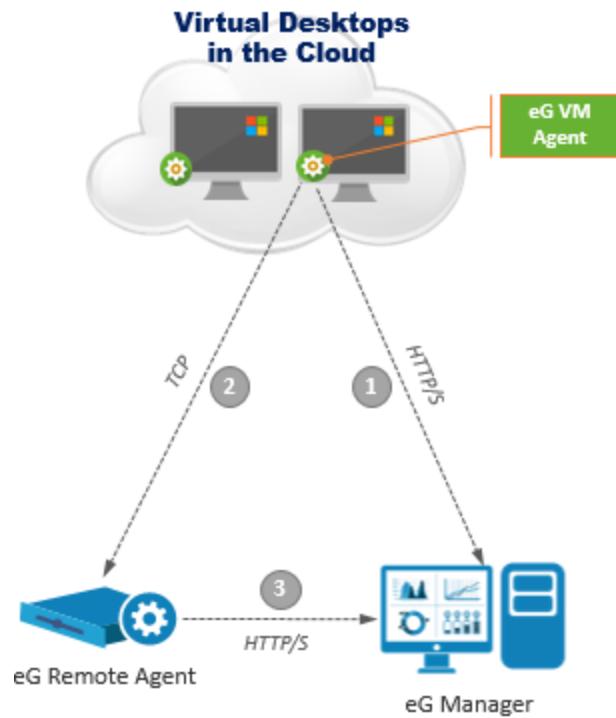


Figure 2.2: The eG VM Agent - eG Remote Agent communication when monitoring desktops on the cloud

- When monitoring Cloud Desktops, the remote agent needs to be in the same subnet as that of the eG VM Agent or accessible on the local network of the eG VM Agent.
- Ensure that the eG VM Agent communicates with the IP address of the remote agent directly and that the communication is not via a NAT or port forward.

## 2.2 Licensing

Virtual desktop monitoring in the cloud is licensed by the number of **concurrent/named users**; and **NOT by servers**.

## Chapter 3: How to Monitor Cloud Desktops Using eG Enterprise?

This involves the following broad steps:

1. Install an eG agent on a remote Windows/Linux host in the environment;
2. Install and start an eG VM Agent on every Windows desktop you want monitored on the cloud;
3. Configure every eG VM Agent with the details of the eG manager it should talk to, so that it can download from the manager the details of the Cloud Desktops component it should monitor, and the remote agent it should report to;
4. Enable a TCP port on the remote agent, using which the VM agent will send metrics to the remote agent;
5. Manage the Cloud Desktops component using the eG administrative interface;
6. View the real-time state and metrics of the managed Cloud Desktops component in the eG monitoring console

To know how to install an eG agent on a Windows/Linux host, refer to the eG Installation Guide. Each of the other steps outlined above (i.e., steps 2-6) has been elaborately discussed in the sub-sections that will follow:

### 3.1 Installing the eG VM Agent

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

- Manually install the eG VM Agent on every Windows virtual desktop on the cloud 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 desktops on the cloud; this way, the eG VM Agent is automatically available in all the virtual desktops that are created using the template;
- Use a software distribution solution such as Microsoft System Center to distribute the eG VM Agent software to Windows desktops on the cloud from a central location;
- Connect to each Windows desktop on the cloud and silently install the eG VM Agent on it, without using the executable that eG Enterprise provides.

The first, third, and fourth installation options alone are discussed here.

### 3.1.1 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 Windows virtual desktop on the cloud, double-click on the **eGVMAgent.exe**, and to install the same on a 64-bit VM, double-click the **eGVMAgent\_64.exe**.
2. Figure 3.1 then appears. Click on the **Next** button in Figure 3.1 to continue.

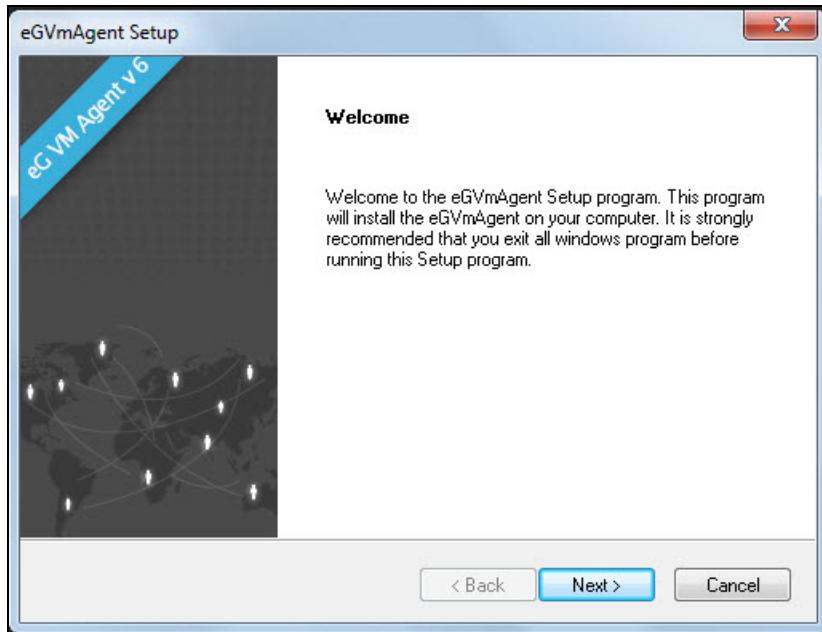


Figure 3.1: Welcome screen of the eG VM Agent installation wizard

3. When Figure 3.2 appears, click on **Yes** to accept the displayed license agreement.

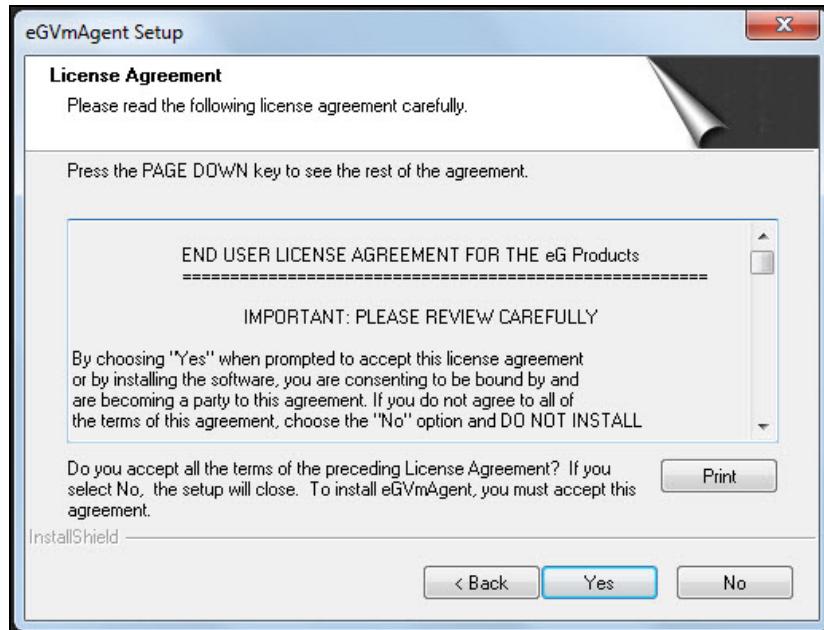


Figure 3.2: Accepting the license agreement

4. Use the **Browse** button in Figure 3.3 to indicate the location in which the agent should be installed, and click the **Next** button to proceed.

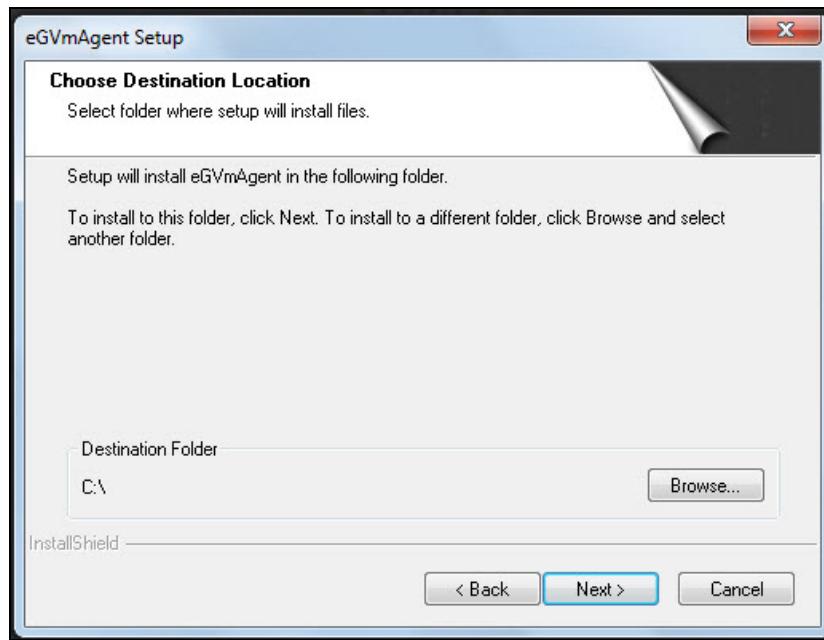


Figure 3.3: Specifying the install directory of the eG VM Agent

5. Next, specify the port through which the VM agent communicates with the eG agent. The

default port is 60001. After port specification, click on the **Next** button in Figure 3.4 to proceed.

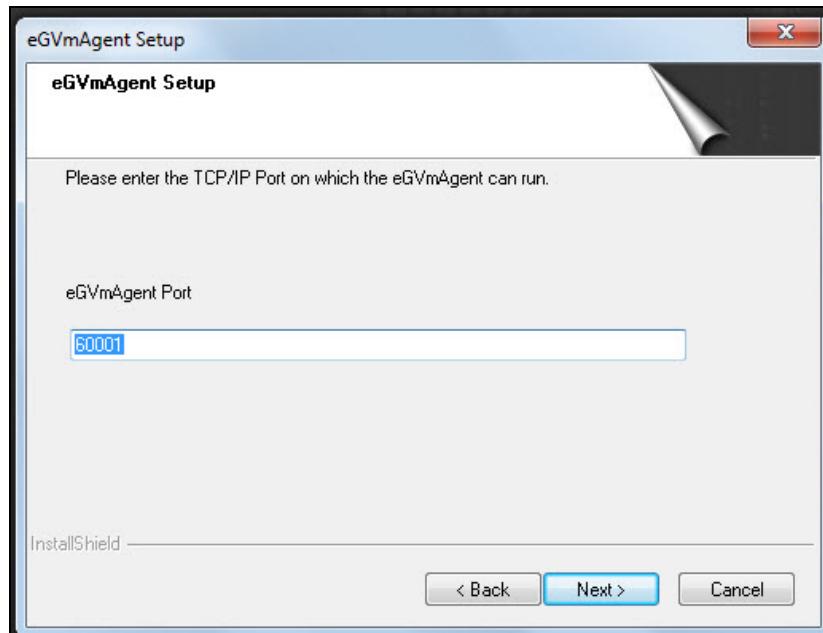


Figure 3.4: Specifying the VM agent port

6. A summary of your specifications then follows (see Figure 3.5). Click **Next** to proceed.

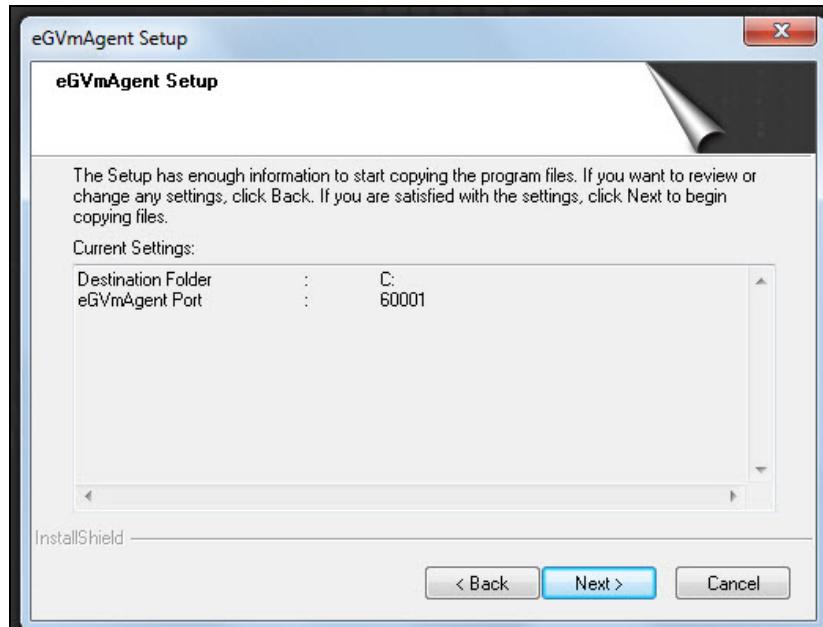


Figure 3.5: A summary of your specifications

7. Finally, click the **Finish** button in Figure 3.6 to complete the installation.

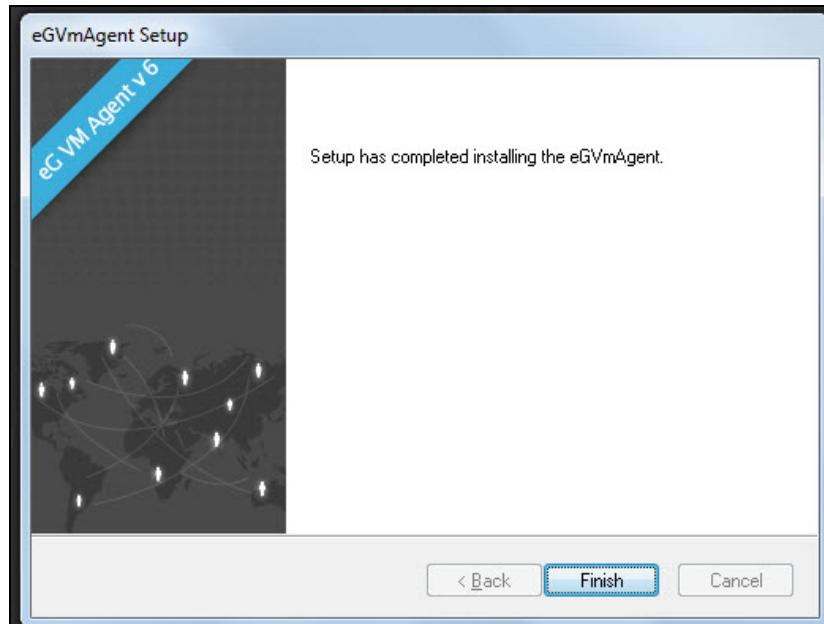


Figure 3.6: Finishing the installation

### 3.1.2 Using SCCM

System Center Configuration Manager (SCCM) is a systems management software product developed by Microsoft for managing large groups of computers running Windows NT, Windows Embedded, OS X, Linux or UNIX, as well as Windows Phone, Symbian, iOS and Android mobile operating systems. Configuration Manager provides remote control, patch management, software distribution, operating system deployment, network access protection and hardware and software inventory.

Like many other software, the eG VM Agent software can also be distributed or remotely deployed on multiple Windows virtual desktops on the cloud using the SCCM.

The broad steps to achieve this are as follows:

1. Building a package for eG VM Agent installation using SCCM
2. Deploying the package using SCCM

Each of these steps are explained in detail below.

#### 3.1.2.1 Building a Package Using SCCM

The steps are as follows:

1. Create a shared folder in your network and allow all Windows virtual desktops on which the eG VM agent is to be installed to access that shared folder.
2. Install the eG VM Agent on any Windows virtual desktop on the cloud.
3. After successful installation of the VM agent, copy the <EG\_VM\_AGENT\_INSTALL\_DIR> - by default, this will be the **eGVmAgent** directory - from that Windows virtual desktop to the shared folder you created in step 1 above.
4. Next, login to the SCCM host and open Notepad. Save the file as **setup.bat**. A sample setup.bat script is provided below:

```
echo ##eG VM Agent Installation## > c:\windows\Temp\egagent.log
GOTO STEP1
:STEP1
echo ##### Starting STEP1 ##### >> c:\windows\Temp\egagent.log
mkdir c:\eGVmAgent
GOTO STEP2
:STEP2
echo #####Starting STEP2##### >> c:\windows\Temp\egagent.log
xcopy /S /I /E eGVmAgent c:\eGVmAgent
GOTO STEP3
:STEP3
echo #####Starting STEP3##### >> c:\windows\Temp\egagent.log
start c:\eGVmAgent\lib\CreateVmAgent.bat
GOTO STEP4
:STEP4
echo #####Starting STEP4##### >> c:\windows\Temp\egagent.log
netsh advfirewall firewall add rule name="eGVMagent Port" dir=in action=allow protocol=TCP
localport=60001
GOTO STEP5
:STEP5
echo #####Starting STEP5##### >> c:\windows\Temp\egagent.log
net start egvmagent
:EOF
```

5. This sample script does the following:
  - Creates a directory called eGVMAgent in the C drive of the target Windows virtual desktop.
  - Copies the contents of the eGVMAgent directory (that was copied to the shared folder in step 3 above) to c:\eGVMAgent.
  - Runs the CreateVMAgent.bat file in the lib folder of C:\eGVMAgent to install the VM agent.
  - Opens the TCP port 60001 on Windows firewall on the VM agent host.
  - Starts the eG VM Agent
6. If you choose to write your own batch file for installing and setting up the VM agent, then make sure it performs all the actions described above. However, you can, if you want, deploy the VM agent in a different directory (and not c:\eGVMAgent) on the virtual desktop. In this case, make sure the script reflects the same.
7. Even if you decide to use the sample script as is, you can still configure a different directory as the install directory of the VM agent. In this case, make sure that all the areas highlighted in **Bold** in the sample script above reflects the new directory.
8. Save the file.
9. Next, start the SCCM's Deployment wizard. Figure 1 will then appear, where you have to provide the **Name** of the package being built. For convenience, name the package as **eGVMAgent**.

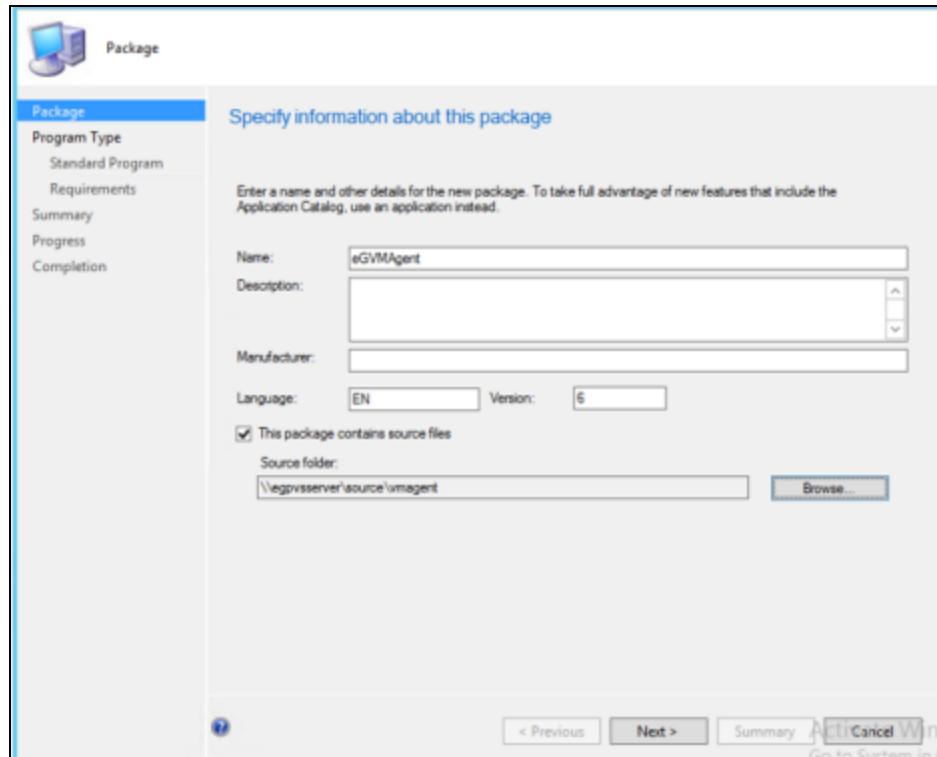


Figure 3.7: Specifying package information

10. Make sure that the **The package contains source files** check box is selected. Then, use the **Browse** button alongside the **Source folder** text box to browse for the shared folder in which the VM agent executable, the \*.iss file, and the setup.bat have been copied.
11. Then, click the **Next** button to proceed. Figure 3.8 will appear.

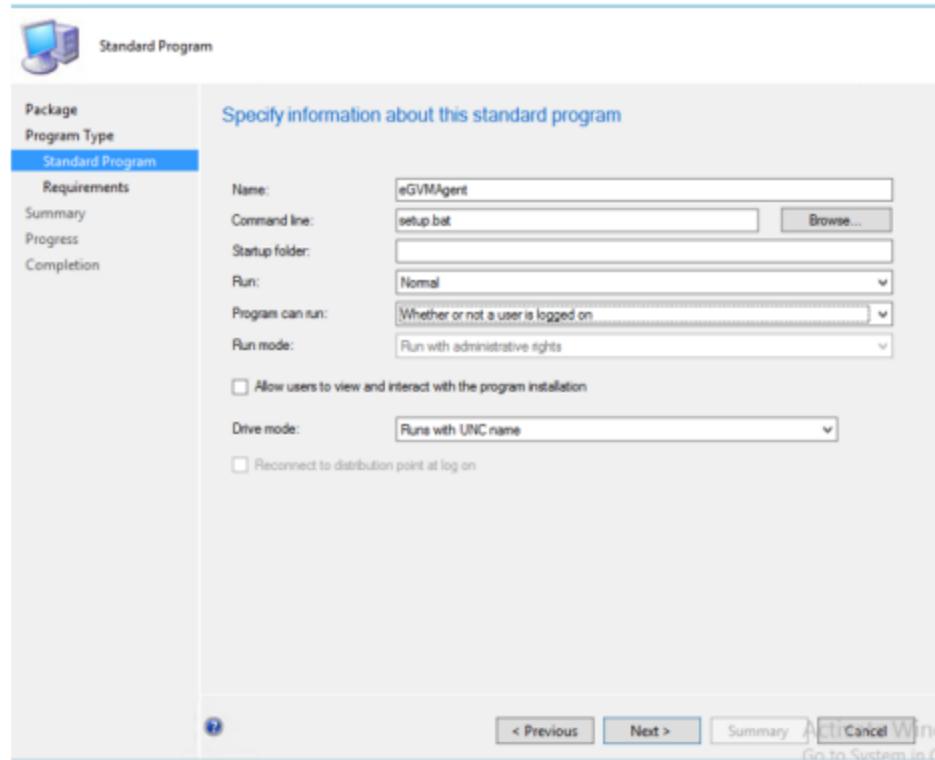


Figure 3.8: Specifying the command to run

12. In Figure 3.8, against **Command line**, enter *setup.bat* - the name of the batch file that contains the commands to deploy the VM agent on a target Windows virtual desktop.
13. Then, from the **Program can run** drop-down, select the **Whether or not a user is logged on** option, so that VM agent installation can proceed whether/not a user is logged into a Windows virtual desktop.
14. Then, click **Next** to move to the next step. Figure 3.9 will appear.

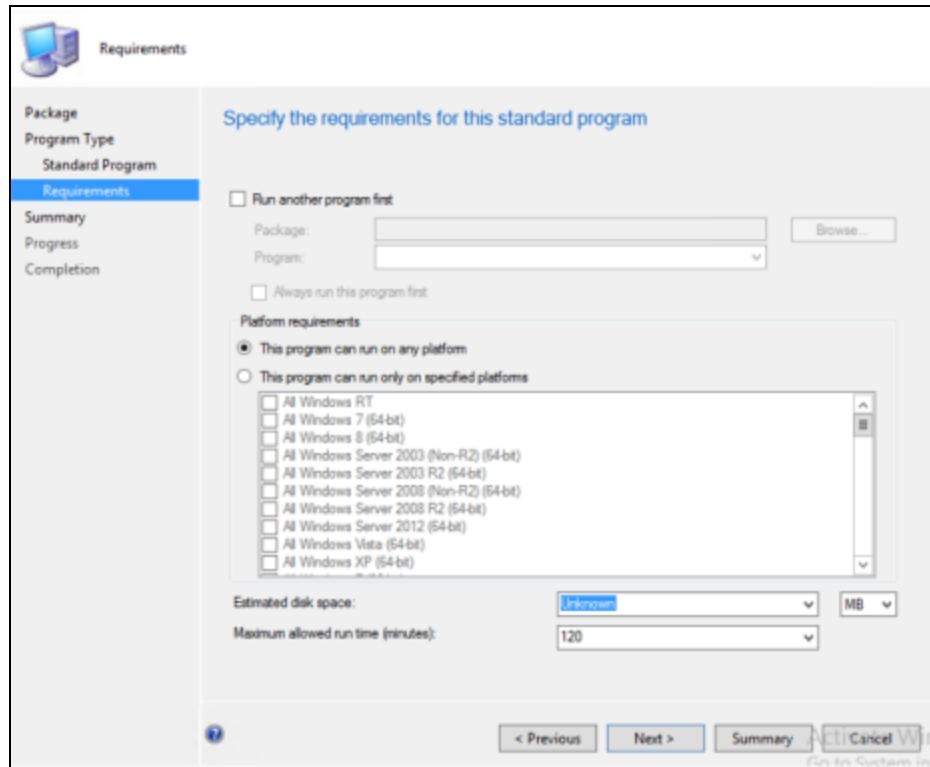


Figure 3.9: Selecting the platform on which the package should be deployed

15. From the **Platform requirements** section of Figure 3.9, select the **This program can run on any platform** option. From the list of platforms, select the Windows platforms on which the package should be deployed. Then, click the **Next** button.
16. Figure 3.10 will then appear displaying a quick summary of the package settings you configured. Review the settings and click the **Next** button to trigger package creation.

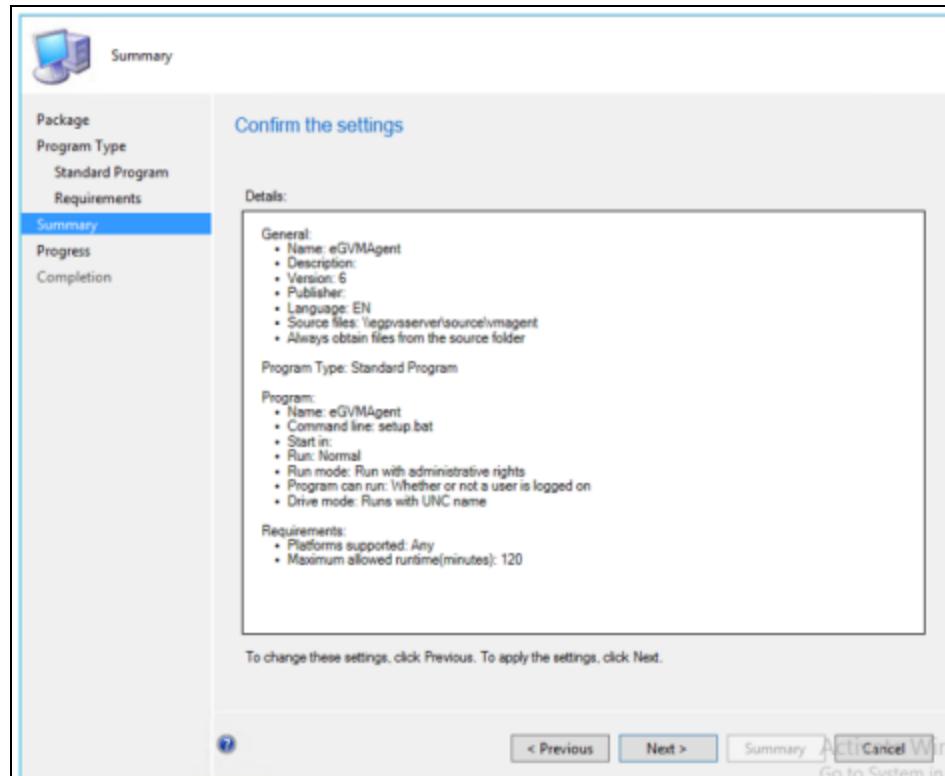


Figure 3.10: A summary of the package specifications

17. If package creation is successful, then Figure 3.11 will then appear informing you of the same.

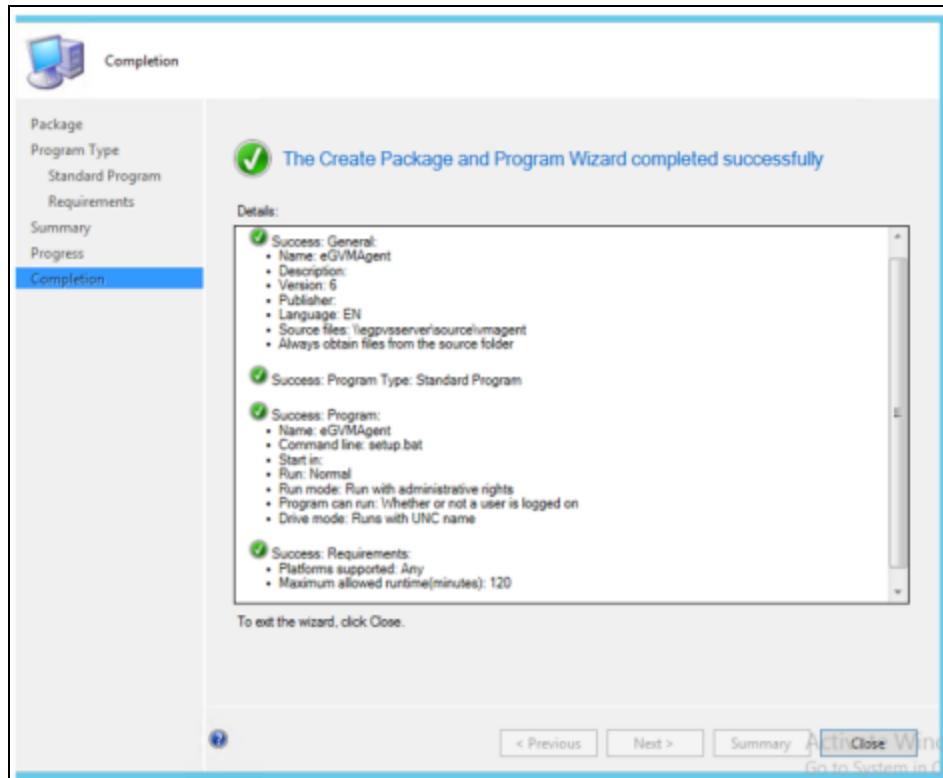


Figure 3.11: The message that appears upon successful package creation

18. Finally, click the **Close** button to exit the wizard.

### 3.1.2.2 Deploying the eGVMAgent Package Using SCCM

Once the package is created, you need to deploy the package on the target Windows virtual desktops on the cloud using SCCM. To do this, follow the steps below:

1. In the **Configuration Manager** console, click **Software Library**.
2. In the **Software Library** workspace, expand **Application Management**, and then click **Packages**.

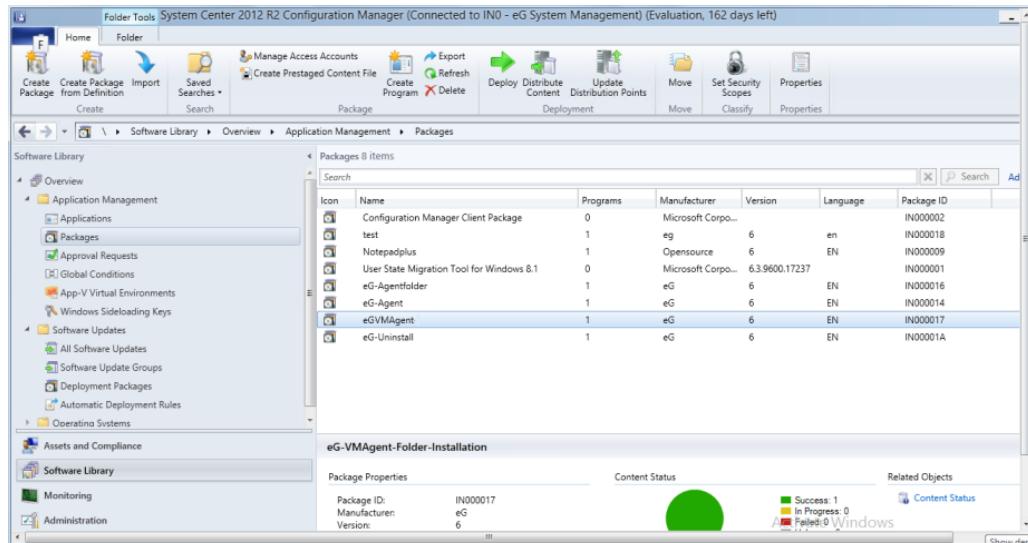


Figure 3.12: Selecting the eGVMAgent package for deployment

3. Select the package that you want to deploy. In our case, this will be the 'eGVMAgent' package you created as described in Section 1.1 above, Then, in the **Home** tab in the **Deployment** group, click **Deploy** (as indicated by the figure below).

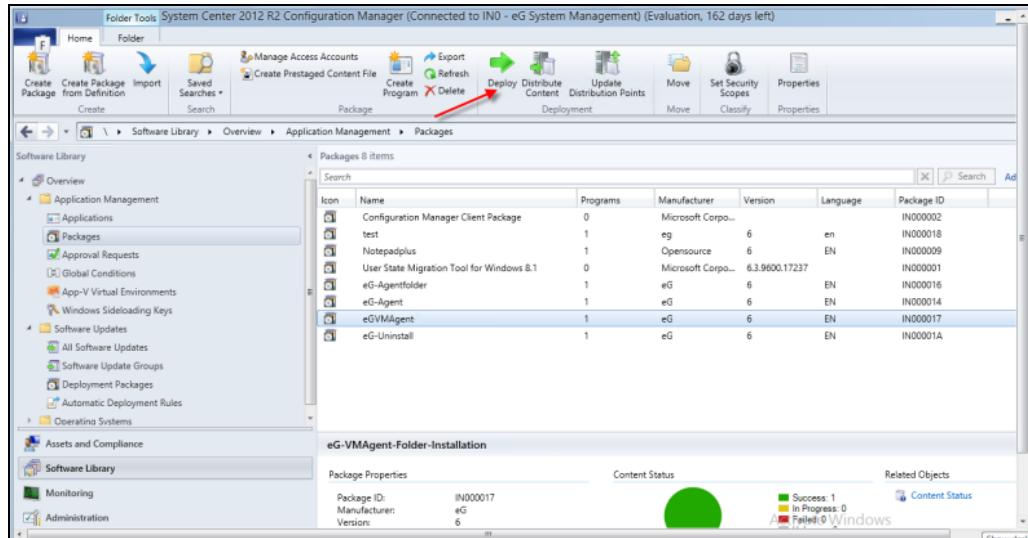


Figure 3.13: Deploying the chosen package

4. On the **General** page of the **Deploy Software Wizard**, specify the name of the package and program that you want to deploy, the collection to which you want to deploy the package and program, and optional comments for the deployment.

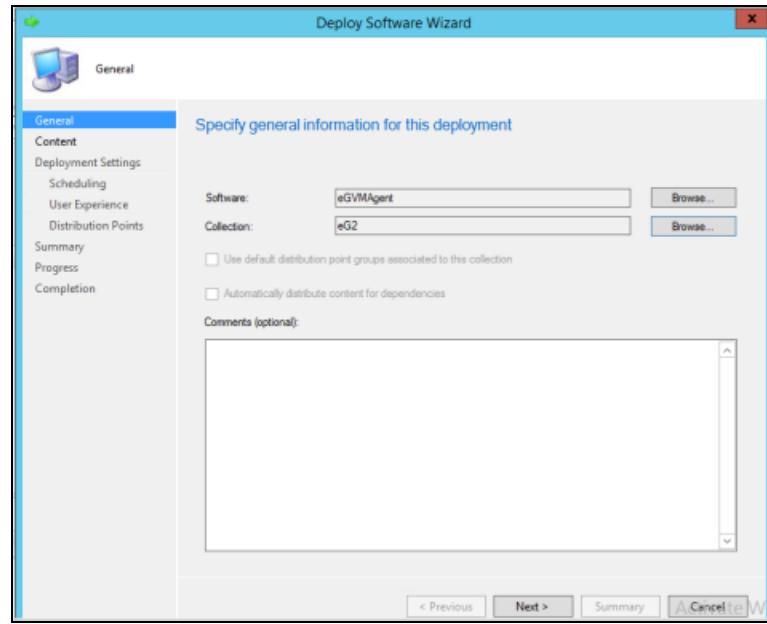


Figure 3.14: Selecting the collection on which the package is to be deployed

5. Select **Use default distribution point groups associated to this collection** if you want to store the package content on the collections default distribution point group. If you did not associate the selected collection with a distribution point group, this option will be unavailable. Click **Next** to proceed.
6. On the **Content** page of the Wizard, click **Add**, and then select the distribution points or distribution point groups to which you want to deploy the content that is associated with this package and program. Click **Next** to proceed.

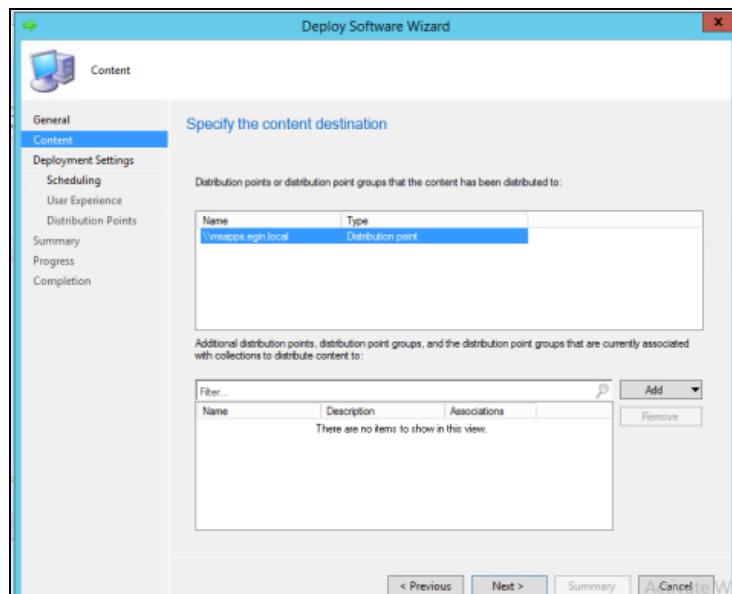


Figure 3.15: Selecting the distribution points

7. On the **Deployment Settings** page of the Wizard, choose **Action** as **install**, pick a purpose for this deployment, and specify options for wake-up packets, and metered connections:

- **Purpose** - Choose from:
  - **Available** - If the application is deployed to a user, the user sees the published package and program in the Application Catalog and can request it on demand. If the package and program is deployed to a device, the user will see it in Software Center and can install it on demand.
  - **Required** - The package and program is deployed automatically, according to the configured schedule. However, a user can track the package and program deployment status and install it before the deadline by using Software Center.
- **Send wake-up packets** – If the deployment purpose is set to **Required** and this option is selected, a wake-up packet will be sent to computers before the deployment is installed to wake the computer from sleep at the installation deadline time. Before you can use this option, computers must be configured for **Wake On LAN**.

Select **Allow clients on a metered Internet connection to download content after the installation deadline, which might incur additional costs** if required.

**Note:**

The **Pre-deploy software to the user's primary device** option is not available when you deploy a package and program.

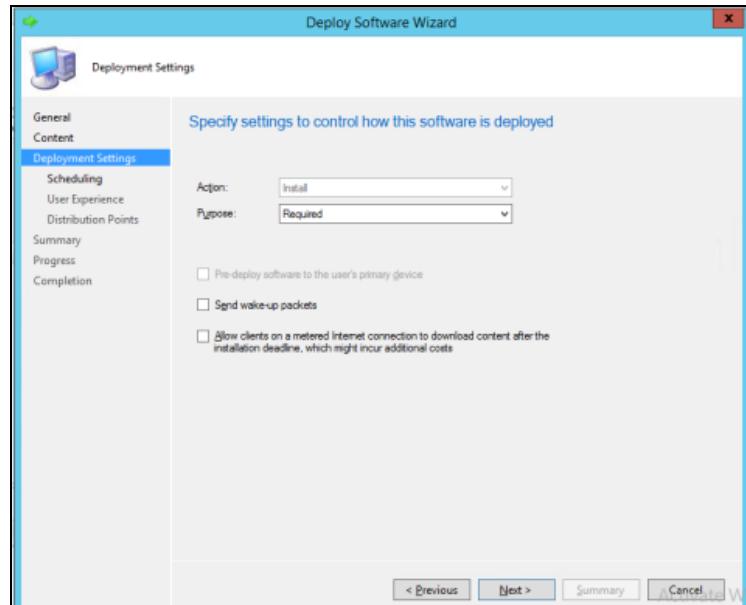


Figure 3.16: Specifying an action and purpose

8. Click **Next** to proceed. On the **Scheduling** page of the Wizard, configure when this package and program will be deployed or made available to client devices.

The options on this page will vary depending on whether the deployment action is set to **Available** or **Required**. Click **Next** to proceed.

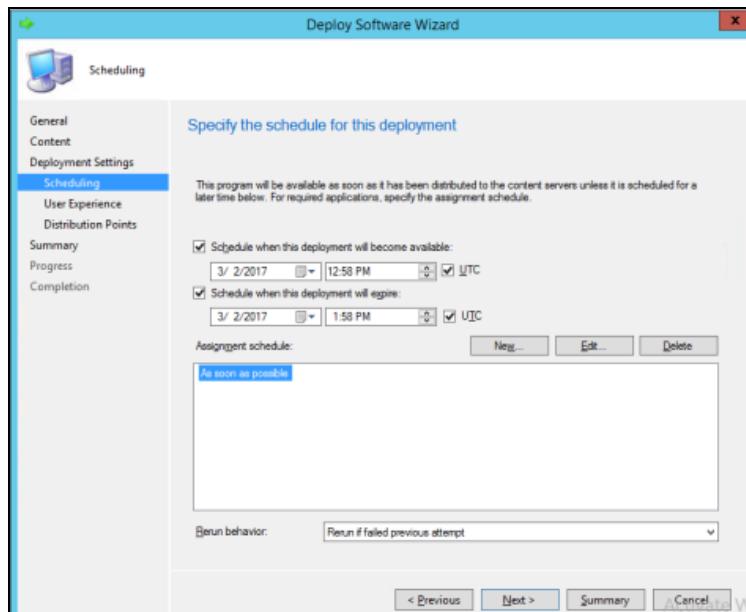


Figure 3.17: Scheduling the package deployment

9. If the deployment purpose is set to **Required**, configure the rerun behavior for the program from the Rerun behavior drop-down list. Choose from the following options:

Rerun behavior	More information
Never rerun deployed program	The program will not be rerun on the client, even if the program originally failed, or the program files are changed.
Always rerun program	The program will always be rerun on the client when the deployment is scheduled, even if the program has already successfully run. This can be useful when you use recurring deployments in which the program is updated, for example with antivirus software.
Rerun if failed previous attempt	The program will be rerun when the deployment is scheduled only if it failed on the previous run attempt.
Rerun if succeeded on previous attempt	The program will be rerun only if it previously ran successfully on the client. This is useful when you use recurring advertisements in which the program is routinely updated, and in which each update requires the previous update to be successfully installed.

10. On the **User Experience** page of the Wizard, specify the following information and then click **Next** to proceed.:

- **Allow users to run the program independently of assignments** – If enabled, users can install this software from the application catalogue regardless of any scheduled installation time.
- **Software installation** – Allows the software to be installed outside of any configured maintenance windows.
- **System restart (if required to complete the installation)** – If the software installation requires a device restart to complete, allow this to happen outside of any configured maintenance windows.
- **Embedded Devices** - For Configuration Manager SP1 only. When you deploy packages and programs to Windows Embedded devices that are write filter enabled, you can specify to install the packages and programs on the temporary overlay and commit changes later, or commit the changes at the installation deadline or during a maintenance window. When you commit changes at the installation deadline or during a maintenance window, a restart is required and the changes persist on the device.

**Note:**

When you deploy a package or program to a Windows Embedded device, make sure that the device is a member of a collection that has a configured maintenance window.

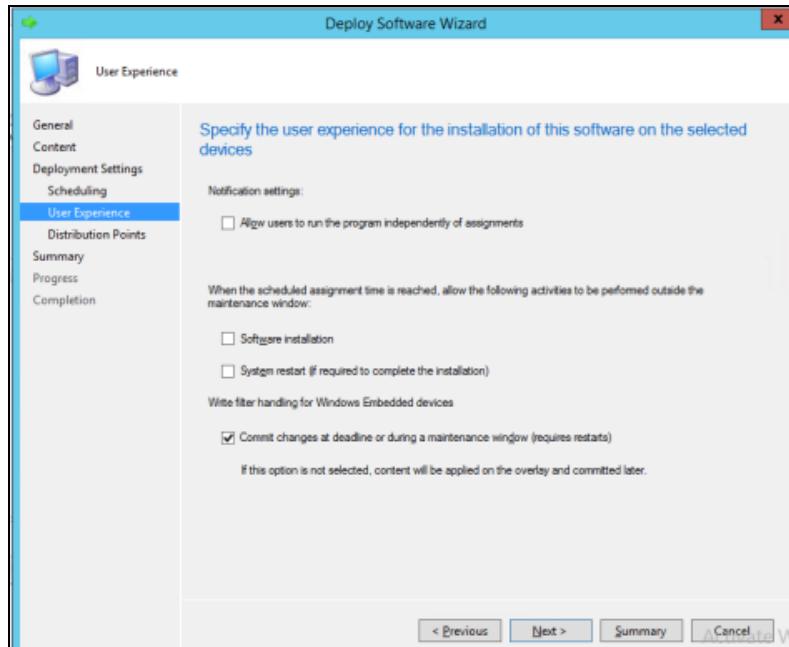


Figure 3.18: Specifying the user experience

11. On the **Distribution Points** page of the Wizard, specify the following information and click **Next** to proceed.:

- **Deployment options** – Specify the actions that a client should take to run program content. You can specify behavior when the client is in a fast network boundary, or a slow or unreliable network boundary.
- **Allow clients to share content with other clients on the same subnet** – Select this option to reduce load on the network by allowing clients to download content from other clients on the network that already downloaded and cached the content. This option utilizes Windows Branch Cache and can be used on computers that run Windows Vista SP2 and later.
- **Allow clients to use a fallback source location for content** – If enabled, clients can search other distribution points in the hierarchy for required content if this is not available on

the specified distribution point or distribution point groups.

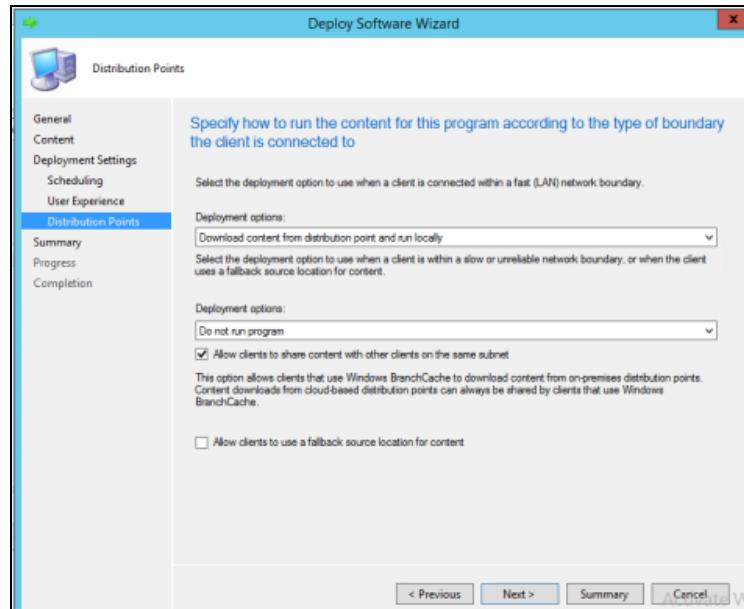


Figure 3.19: Specifying how to run the content for the program

12. On the **Summary** page of the Wizard, review the actions that will be taken and then click **Next** to proceed.

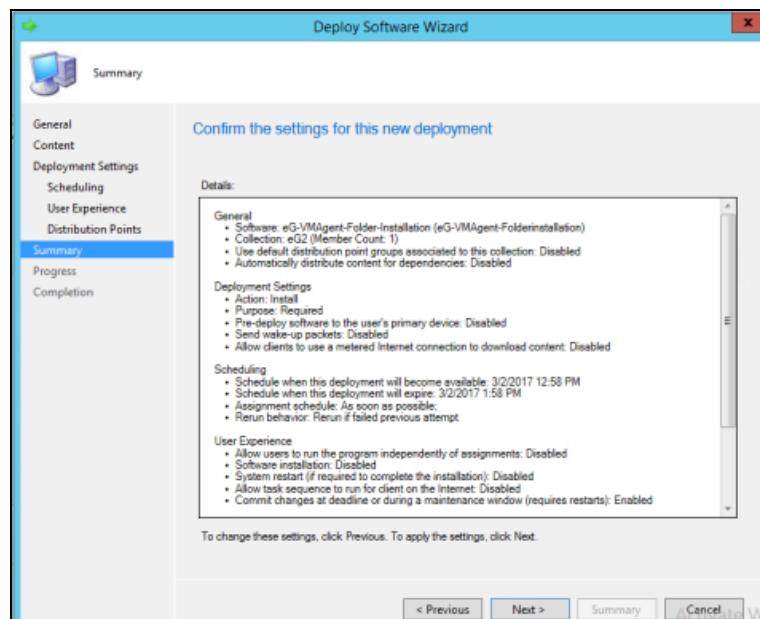


Figure 3.20: Viewing a summary of the settings

13. Finally, click **Close** to exit the wizard.

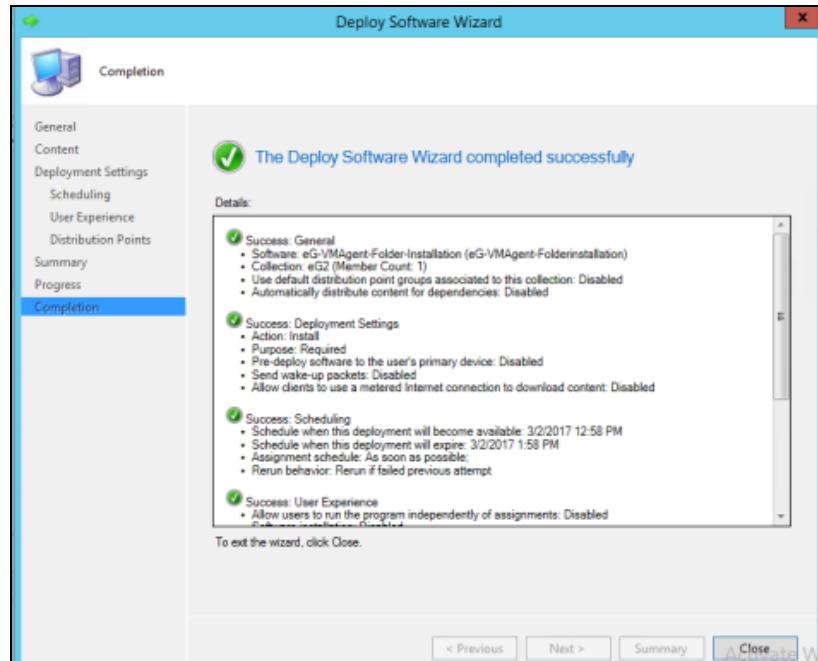


Figure 3.21: Completion of deployment wizard

14. You can view the deployment in the **Deployments** node of the **Monitoring** workspace and in the details pane of the package deployment tab when you select the deployment.

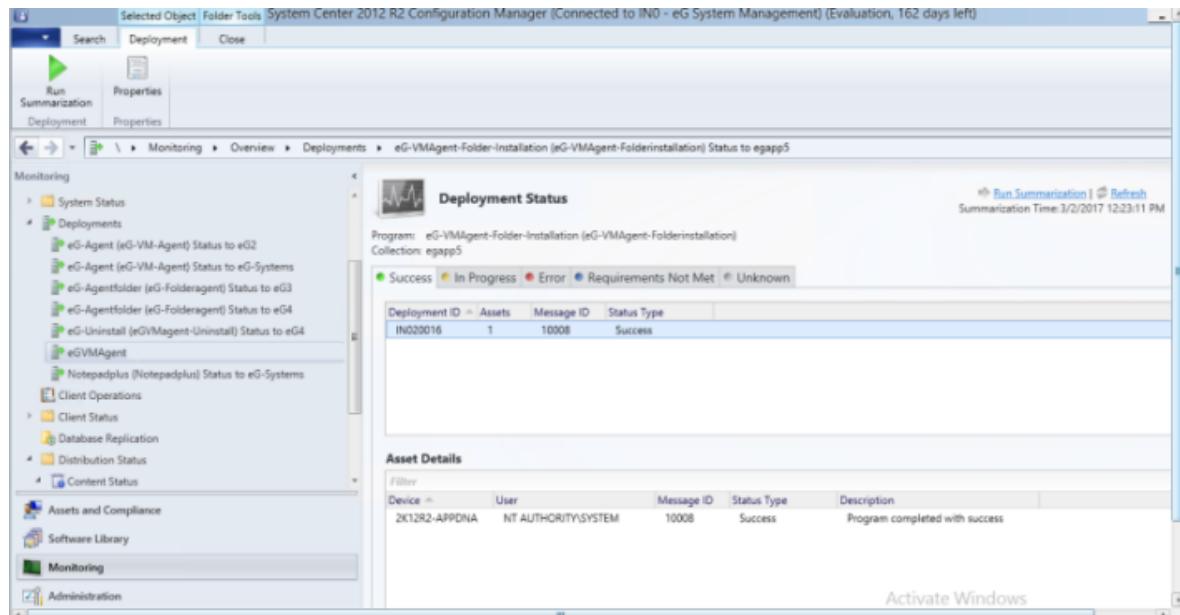


Figure 3.22: Viewing the status of the deployment

**Note:**

If you configured the option **Run program from distribution point** on the **Distribution Points** page of the Deploy Software Wizard, do not clear the option **Copy the content in this package to a package share on distribution points**, because this will make the package unavailable to run from distribution points.

### 3.1.3 Silent Installation of the eG VM Agent

To silently install the eG VM agent on Windows virtual desktops on the cloud, 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.

#### 3.1.3.1 Creating a Silent Mode Script

For this, follow the procedure detailed below:

1. Login to a target Windows desktop on the cloud.
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\leGVMAgent.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.

#### 3.1.3.2 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 virtual desktop on the cloud, where the script file containing the inputs for installation resides.
2. Copy the script file from this desktop to the Windows virtual desktop 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 virtual desktop.
4. Next, on the target Windows virtual desktop, 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\leGVMAgent.iss"**

5. Upon successful execution, this command will automatically install the eG VM Agent on the target.
6. You can then repeat steps 1-5 on each Windows virtual desktop on the cloud where you want to install the eG VM Agent.

## 3.2 Configuring the eG VM Agent to Communicate with the eG Manager

The eG VM Agent, once installed and started, will begin collecting metrics on internal performance of the target Windows virtual desktop. However, the VM agent has no knowledge of the remote agent to which it needs to send the metrics; nor does it know to which Cloud Desktop component these metrics pertain. Without this knowledge, the eG VM Agent will not be able to transmit the right 'inside view' metrics to the right eG remote agent.

For this purpose, the eG VM agent has to be configured with the nick name of a Cloud Desktop component that has been managed in the eG manager. Also, to determine the remote agent that has been assigned to this Cloud Desktop component, the VM agent has to be configured to communicate with the eG manager and download the details of that Cloud Desktop component.

To achieve this, follow the steps below:

1. Login to the Windows virtual desktop hosting the eG VM agent.
2. Go to the command prompt and switch to the <EGVMAGENT\_INSTALL\_DIR>\lib directory.

3. Run the following command:

**SetManagerInfo.bat**

4. The script will first request you to enter the nick name of the Cloud Desktop component for which this VM agent collects metrics. Make sure that this name matches the nick name of the Cloud Desktop component that has been managed in the eG manager.

```
This script will configure the eG Manager to which the VM agent will connect to.
```

```
Please enter the Nickname of this VM Agent: eGDevCloudAcc
```

```
This name should match the nickname of the Cloud Desktop component
```

5. Next, specify the IP address and then the port number of the eG manager with which this VM agent should communicate.

```
Please enter the eG Manager IP/Name to which the VM Agent should connect:  
192.168.8.149
```

```
Please enter the Port on which the VM Agent connects to the eG Manager: 7077
```

6. Then, indicate whether/not the eG manager is SSL-enabled. If it is, then specify **Yes** at this prompt. If not, enter **No**.

```
Is the eG Manager SSL enabled (Yes/No) ? No
```

7. If the VM agent should use a proxy server to communicate with the eG manager, then specify Yes at the next prompt. If not, then enter No.

```
Should the VM Agent use a proxy server to connect to the eG Manager (Yes/No) ? No
```

8. If you entered Yes at step 7, then you will be required to provide the IP and port number of the proxy server via which the VM agent and manager should communicate.

```
Please enter the proxy IP/Name: 10.10.1.13
```

```
Please enter the proxy port: 6789
```

9. Then, specify Yes at the next prompt if the proxy server to be used requires authentication. If the proxy server does not require authentication, enter No.

```
Does the proxy require user authentication (Yes/No) ? Yes
```

10. If you had specified Yes at step 9, then you will be required to mention the proxy user name and password for authentication.

```
Please enter the proxy username: egtest
```

```
Please enter the proxy password: w3lc0m3
```

- Once the details are provided, the script will then try to connect to the eG manager specified at step 5, download the details of the managed Cloud Desktop component from the eG manager, and write it to a eg\_managersetup.ini file on the VM agent host. Then, the script will restart the eG VM Agent service.

```
Trying to connect to the eG Manager...
```

```
Connected to http://192.168.8.149:7077
```

```
Details of the eG Manager have been successfully configured in the file:  
C:\eGVmAgent\lib\eg_managersetup.ini
```

```
The eGVmAgent service is stopping.
```

```
The eGVmAgent service was stopped successfully.
```

```
The eGVmAgent service is starting.
```

```
The eGVmAgent service was started successfully.
```

### 3.3 Enabling a TCP Port on the eG Remote Agent

By default, the remote agent, like any other eG agent, does not listen on any port. However, to be able to receive metrics from the eG VM Agent on a cloud desktop, the remote agent has to be configured with a TCP port. For this, do the following:

1. Login to the eG manager host.
2. Edit the **eg\_tests.ini** file (in the <EG\_MANAGER\_INSTALL\_DIR>\manager\config directory, on Windows; on Unix, this file will be in the /opt/egurkha/manager/config directory).
3. To enable the remote agent to listen on a port, set the **RemoteAgentPortEnabled** flag in the **[AGENT\_SETTINGS]** section of the file to **yes**.
4. To configure the port at which the agent should listen, specify a port number against the **RemoteAgentPort** parameter in the **[AGENT\_SETTINGS]** section of the file. By default, this parameter will be set to **60010**. If you want the agent to listen on a different port, change this port number.
5. Finally, save the file.
6. If the remote agent is behind a firewall, then make sure that you configure the firewall to allow one-way communication between the eG VM agent and the remote agent, via the port specified against **RemoteAgentPort**.

## 3.4 Managing a Cloud Desktops Component in eG Enterprise

Follow the procedure detailed below to manage a Cloud Desktops component in eG Enterprise:

1. Login to the eG administrative interface as a user who has been assigned the Admin role.
2. Invoke the Admin tile menu and follow the Infrastructure -> Components -> Add/Modify menu sequence in it.
3. Select *Cloud Desktops* as the **Component type** from the page that appears next. Then, click the **Add New Component** button alongside.
4. Figure 3.23 will then appear.

The screenshot shows a configuration dialog for adding a new component. At the top, there are dropdown menus for 'Category' (set to 'All') and 'Component type' (set to 'Cloud Desktops'). Below these are two sections: 'Component information' and 'Monitoring approach'. In the 'Component information' section, the 'Nick name' field is populated with 'CWGCloudDesktop'. In the 'Monitoring approach' section, the 'Agentless' checkbox is checked. The 'OS' dropdown is set to 'Other'. The 'Mode' dropdown is also set to 'Other'. The 'Remote agent' dropdown is set to 'eGLAP0043-PC', and the 'External agents' dropdown is also set to 'eGLAP0043-PC'. At the bottom of the dialog is a large 'Add' button.

Figure 3.23: Adding a Cloud Desktops component

5. In Figure 3.23, specify the following:

- Provide a **Nick name** for the Cloud Desktops component being added. Make sure that the nick name you specify here is the same as the nick name you provided for the Cloud Desktop component when configuring the eG VM Agent (see Section [Section 3.2](#) ).
- Select the **Agentless** check box, set the **Mode** to *Other*, and pick the **Remote agent** that should monitor the Cloud Desktops component. Ensure that the remote agent you choose here is the same as the remote agent you specified when configuring the eG VM agent (see Section [Section 3.2](#) ).

- Assign an **External agent** to the Cloud Desktops component, and click the **Add** button to add the component into the eG Enterprise system.

## Chapter 4: Monitoring the Cloud Desktops Component

Log into the eG monitoring interface for viewing the current state of the managed Cloud Desktops component and the performance statistics it reports.

In the eG monitoring console, eG Enterprise uses a specialized *Cloud Desktops* monitoring model to represent the real-time state of the Windows cloud desktops.

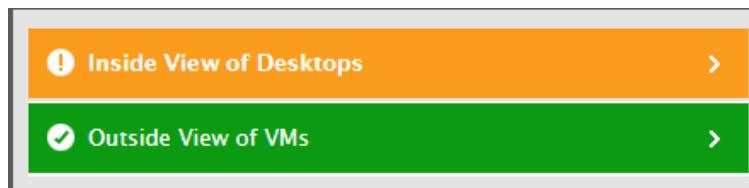


Figure 4.1: Layer model of the Cloud Desktop component

The tests mapped to these layers report metrics that have been collected from those VM agents that have been configured to monitor the same Cloud Desktop component - i.e., that have been configured with the same Cloud Desktop component nick name.

These tests track user logins and logouts on the Windows virtual desktop on the cloud, monitor user activities, and report on the resource impact of those activities on the internal health of that desktop. Using the measures reported by these tests, administrators can find quick and accurate answers to the following performance queries:

- How many desktops are powered on simultaneously on the cloud? Which ones are they and which OS are they running on?
- Which users are logged on and when did each user login?
- How much CPU, memory, GPU, and network resources is each desktop taking?
- What is the typical duration of a user session?
- Who has the peak usage times?
- What applications are running on each desktop? Which applications are the top CPU/memory consumers on the desktop?
- Are any users connecting to the desktop via Blast / PCoIP? If so, is any such user's experience with the desktop sub-par? What could be impacting user experience?

The sub-sections that follow will discuss each layer of Figure 1 in great detail.

## 4.1 The Outside View of VMs Layer

The tests mapped to this layer automatically discover the IP address of every Windows virtual desktop where a VM agent configured with this Cloud Desktop's nick name is running. The powered-on state of each desktop is reported. Also, the user logins and logouts on each virtual desktop is tracked, and the recent logins and logouts are reported.



Figure 4.2: The tests mapped to the Outside View of VMs layer

### 4.1.1 VM Logins - Cloud Test

This test monitors the user logins to the cloud-hosted Windows desktops and reports the total count of logins and logouts.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for the 'Cloud Desktops' component being monitored

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the

Parameter	Description
	<i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to <i>1:1</i> , it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
Detailed Diagnosis	<p>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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current sessions	Indicates the number of user sessions that are currently active across all the Windows virtual desktops that are reporting metrics for the target Cloud Desktops component	Number	This is a good indicator of the session load on the Windows virtual desktops.
New logins	Indicates the number of new logins to the Windows virtual desktops	Number	<p>A consistent zero value could indicate a connection issue.</p> <p>If this measure reports a non-zero value, use the detailed diagnosis of the measure to know which user logged into which virtual desktop when.</p>

Measurement	Description	Measurement Unit	Interpretation
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.

Use the detailed diagnosis of the *New logins* measure to know which user logged into which virtual desktop when.

Details of current user sessions		
VM NAME	USERNAME	LOGIN TIME
Apr 23, 2018 18:25:50		
VMWHv72-NP1	citrix\ctxuser	03/02/2018 16:43:18
DESKTOP-6zsd2ajg8	citrix\eguser	06/03/2018 10:20:33
DESKTOP-91T8MCQ	citrix\alan	04/03/2018 17:13:32
VMWHv70-NP1	citrix\james	12/01/2018 11:53:06
VMWHv71-NP1	citrix\tommy	03/02/2018 22:04:18
VMWHv70-NP2	citrix\willie	26/01/2018 13:46:22
VMWHv71-NP2	citrix\jeff	06/03/2018 14:12:18
VMWHv72-NP2	citrix\keith	09/01/2018 23:01:18
DESKTOP-6zd1ci8o8	citrix\kevin	05/02/2018 21:43:18
DESKTOP-91Tu7r34hv8Q	citrix\sandy	19/02/2018 08:43:18
DESKTOP-b36gsz7lg8	citrix\mike	17/02/2018 06:43:18
DESKTOP-91T8jahatat	citrix\scott	03/02/2018 16:43:18
DESKTOP-6zsd3ajg8	citrix\steven	06/03/2018 10:20:33
DESKTOP-91T6MCQ	citrix\johnny	26/01/2018 13:46:22
DESKTOP-6zsd2jwV10	citrix\timothy	04/03/2018 17:13:32

Figure 4.3: The detailed diagnosis of the New logins measure

#### 4.1.2 Virtual Machines - Cloud Test

This test reports the number and names of Windows virtual desktops (on the cloud) for which the VM agents reported metrics during the last hour. Additionally, the test also reveals those VMs that are currently powered-on, and those VMs that have users logged into them presently. This way, administrators can quickly identify the powered-off VMs and those that are not in use currently.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for the 'Cloud Desktops' component being monitored

## Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to <b>1:1</b> , it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total VMs reported in the last one hour	Indicates the number of Windows virtual desktops (on the cloud) that reported metrics during the last hour.	Number	To know which virtual desktops reported metrics, use the detailed diagnosis of this measure.
Currently powered on VMs	Indicates the number of VMs that are powered-on currently.	Number	If the value of this measure is the same as that of the <i>Total VMs reported in the last one hour</i> measure, it means that all Windows virtual desktops are presently in a powered-on state. On the other hand, if the value of this measure is lesser than the value of the <i>Total VMs reported in the last one hour</i> measure, it implies that some virtual desktops are powered-off currently. The value 0 for this measure indicates that none of the VMs that reported in the last hour are powered-on currently.
VMs with users	Indicates the number of Windows virtual desktops to which users are currently logged in.	Number	Use the detailed diagnosis of this measure to know to which virtual desktops users have logged in currently.
VMs without users	Indicates the number of Windows virtual desktops to which no user is currently logged in.	Number	Use the detailed diagnosis of this measure to know to which virtual desktops no user is presently logged in.

Using the detailed diagnosis of the *Total VMs reported in the last hour* measure to you can learn the name, IP address, and the OS of the reporting VMs.

Details of Cloud VMs			
VM NAME	IP ADDRESS	OS	
Apr 23, 2018 18:25:00			
DESKTOP-91T8MCQ	192.168.9.148,fe80::2c3e:db66:70e5:7c54	Windows 10 Pro	
DESKTOP-6zsd2qjg8	192.168.8.119	Windows 8 Pro	
VMWHv72-NP1	192.168.9.167,fe80::c086:ff5:89d1:b460	Windows 10 Pro	
VMWHv70-NP1	192.168.10.116,fe80::781e:dcab:2d39:911f	Windows 10 Pro	
VMWHv71-NP1	192.168.9.11,fe80::250:56ff:fea8:703c	Windows 8 Pro	
VMWHv70-NP2	192.168.11.163,fe80::f1bd:44f4:b0b:648	Windows 7	
VMWHv71-NP2	192.168.8.187,fe80::e1d1:7e2:ccb:4416	Windows 7	
VMWHv72-NP2	192.168.9.44,fe80::91f2:18b:dbd8:a161	Windows 10 Pro	
DESKTOP-6zd1ci8o8	192.168.9.197,fe80::8141:a6aa:1785:6fa	Windows 7	
DESKTOP-91Tu7r34hv8Q	192.168.9.234,fe80::b10d:9fc5:154b:5dba	Windows 8 Pro	
DESKTOP-B36gsz7lg8	192.168.8.71,fe80::c1f0:f9ac:341e:2921	Windows 8 Pro	
DESKTOP-91T8jpahtat	192.168.11.227,fe80::18a8:e257:7fed:f3f1	Windows 10 Pro	
DESKTOP-6zsd3qjg8	192.168.9.149	Windows 10 Pro	
DESKTOP-91T6MCQ	192.168.9.170,fe80::dc8:9f2f:4e9e:a306	Windows 7	
DESKTOP-6zsd2jwV10	192.168.9.157	Windows 8 Pro	

Figure 4.4: The detailed diagnosis of the Total VMs reported in the last hour measure

Use the detailed diagnosis of the *VMs with users* measure to know which Windows virtual desktops (on the cloud) have users logged in currently and the name of the user who is logged in.

Details of VMs with users				
VM NAME	IP ADDRESS	OS	USER NAME	
Apr 23, 2018 18:25:00				
VMWHv72-NP1	192.168.9.167,fe80::c086:ff5:89d1:b460	Windows 10 Pro	citrix\ctxuser	
DESKTOP-6zsd2qjg8	192.168.9.170,fe80::dc8:9f2f:4e9e:a306	Windows 8 Pro	citrix\eguser	
DESKTOP-91T8MCQ	192.168.11.49	Windows 10 Pro	citrix\alan	
VMWHv70-NP1	192.168.10.116,fe80::781e:dcab:2d39:911f	Windows 10 Pro	citrix\james	
VMWHv71-NP1	192.168.9.11,fe80::250:56ff:fea8:703c	Windows 8 Pro	citrix\tommy	
VMWHv70-NP2	192.168.11.163,fe80::f1bd:44f4:b0b:648	Windows 7	citrix\willie	
VMWHv71-NP2	192.168.8.187,fe80::e1d1:7e2:ccb:4416	Windows 7	citrix\jeff	
VMWHv72-NP2	192.168.9.44,fe80::91f2:18b:dbd8:a161	Windows 10 Pro	citrix\keith	
DESKTOP-6zd1ci8o8	192.168.9.197,fe80::8141:a6aa:1785:6fa	Windows 7	citrix\kevin	
DESKTOP-91Tu7r34hv8Q	192.168.9.234,fe80::b10d:9fc5:154b:5dba	Windows 8 Pro	citrix\sandy	
DESKTOP-B36gsz7lg8	192.168.8.71,fe80::c1f0:f9ac:341e:2921	Windows 8 Pro	citrix\mike	
DESKTOP-91T8jpahtat	192.168.11.227,fe80::18a8:e257:7fed:f3f1	Windows 10 Pro	citrix\scott	
DESKTOP-6zsd3qjg8	192.168.9.149	Windows 10 Pro	citrix\steven	
DESKTOP-91T6MCQ	192.168.9.170,fe80::dc8:9f2f:4e9e:a306	Windows 7	citrix\johnny	
DESKTOP-6zsd2jwV10	192.168.9.157	Windows 8 Pro	citrix\timothy	

Figure 4.5: The detailed diagnosis of the VMs with users measure

## 4.2 The Inside View of Desktops Layer

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 Windows virtual desktops on the cloud to analyze how well each desktop utilizes the resources that are allocated to it, and how well it handles user sessions, TCP traffic, and network loading.

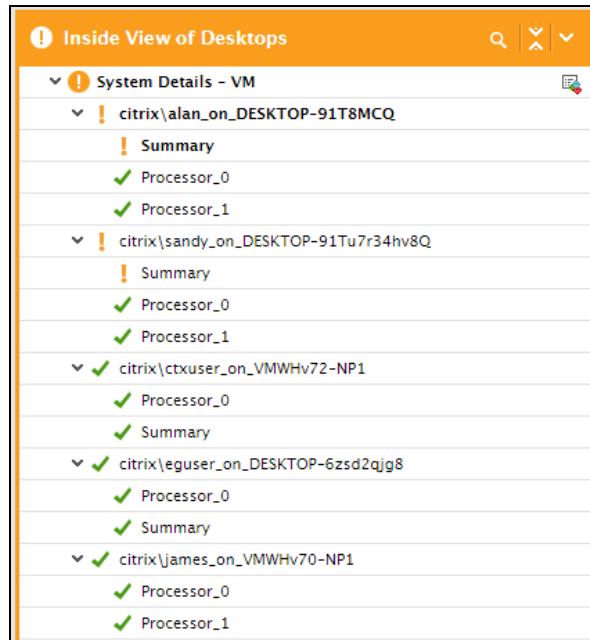


Figure 4.6: The tests mapped to the Inside View of Desktops layer

#### 4.2.1 Disk Activity - VM Test

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

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every combination of *user\_on\_virtualdesktop:disk\_partition*

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a

Parameter	Description
	light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to <i>1:1</i> , it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
Detailed Diagnosis	<p>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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### 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	Indicates the percentage	Percent	

Measurement	Description	Measurement Unit	Interpretation
virtual disk	of elapsed time that the selected disk drive is busy servicing read requests.		
Percent writes to virtual disk	Indicates the percentage of elapsed time that the selected disk drive is busy servicing write requests.	Percent	
Virtual disk read time	Indicates the average time in seconds of a read of data from the disk.	Secs	
Virtual disk write time	Indicates the average time in seconds of a write of data from the disk.	Secs	
Avg. queue for virtual disk	Indicates the average number of both read and write requests that were queued for the selected disk during the sample interval.	Number	
Current queue for virtual disk	The number of requests outstanding on the disk at the time the performance data is collected.	Number	This measure includes requests in service at the time of the snapshot. This is an instantaneous length, not an average over the time interval. Multi-spindle disk devices can have multiple requests active at one time, but other concurrent requests are awaiting service. This counter might reflect a transitory high or low queue length, but if there is a sustained load on the disk drive, it is likely that this will be consistently high. Requests experience delays proportional to the length of this queue minus the number of spindles on the disks. This difference should average less than two for good

Measurement	Description	Measurement Unit	Interpretation
			performance.
Reads from virtual disk	Indicates the number of reads happening on a logical disk per second.	Reads/Sec	A dramatic increase in this value may be indicative of an I/O bottleneck on the guest.
Data reads from virtual disk	Indicates the rate at which bytes are transferred from the disk during read operations.	KB/Sec	A very high value indicates an I/O bottleneck on the guest.
Writes to virtual disk	Indicates the number of writes happening on a local disk per second.	Writes/Sec	A dramatic increase in this value may be indicative of an I/O bottleneck on the guest.
Data writes to virtual disk	Indicates the rate at which bytes are transferred from the disk during write operations.	KB/Sec	A very high value indicates an I/O bottleneck on the guest.
Disk service time	Indicates the average time that this disk took to service each transfer request (i.e., the average I/O operation time)	Secs	A sudden rise in the value of this measure can be attributed to a large amount of information being input or output. A consistent increase however, could indicate an I/O processing bottleneck.
Disk queue time	Indicates the average time that transfer requests waited idly on queue for this disk.	Secs	Ideally, the value of this measure should be low.
Disk IO 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.
Avg IO read size	Indicates the average number of bytes transferred from disk during read operations.	KB	Larger I/Os tend to have higher latency (for example, BACKUP/RESTORE operations issue 1 MB transfers by default).

Measurement	Description	Measurement Unit	Interpretation
Avg I/O write size	Indicates the average number of bytes transferred into disk during write operations.	KB	
Split IO	Reports the rate at which the operating system divides I/O requests to the disk into multiple requests.	Splits/Sec	A split I/O request might occur if the program requests data in a size that is too large to fit into a single request or if the disk is fragmented. Factors that influence the size of an I/O request can include application design, the file system, or drivers. A high rate of split I/O might not, in itself, represent a problem. However, on single-disk systems, a high rate for this counter tends to indicate disk fragmentation.

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_Rate (Ops/sec)
<b>Jan 03, 2008 05:44:08</b>							
	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	csss	3320	3320	7.01	0	0
	2396	vmiprvse	194.82	125.43	1.33	69.39	1.33
<b>Jan 03, 2008 05:33:43</b>							
	1032	svchost	770.92	137.38	3	633.54	2.33
	2396	vmiprvse	194.73	125.37	1.33	69.36	1.33
<b>Jan 03, 2008 05:24:01</b>							
	1032	svchost	770.82	137.36	3	633.46	2.33
	2396	vmiprvse	194.7	125.36	1.33	69.35	1.33
	628	csss	16	16	0.67	0	0
<b>Jan 03, 2008 05:14:10</b>							
	4	System	45069.31	45069.31	86.03	0	0
	1032	svchost	20346.68	7702.94	9.67	12643.73	1.33
	628	cmd	73.69	73.69	1.67	0	0
	2876	tomcat	16.34	0	0	16.34	0.67
	628	csss	16	16	1.33	0	0

Figure 4.7: The detailed diagnosis of the Percent virtual busy measure

#### 4.2.1.1 Configuring Users for VM Monitoring

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

Figure 4.8: Configuring a VM test

Upon clicking, Figure 4.9 will appear, using which the VM user details can be configured.

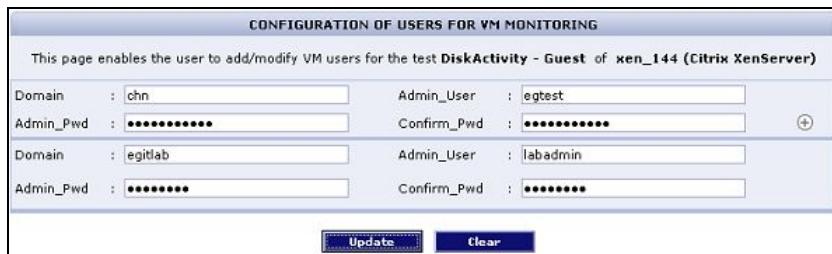
Figure 4.9: The VM user configuration page

To add a user specification, do the following:

1. First, provide the name of the **Domain** to which the VMs belong (see Figure 4.9). 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 Troubleshooting the Failure of the eG Remote Agent to Connect to or Report Measures for Linux Guests 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 4.9. This will allow you to add one more user specification as depicted by Figure 4.10.



CONFIGURATION OF USERS FOR VM MONITORING					
This page enables the user to add/modify VM users for the test <b>DiskActivity - Guest</b> of <b>xen_144 (Citrix XenServer)</b>					
Domain	:	chn	Admin_User	:	egtest
Admin_Pwd	:	[REDACTED]	Confirm_Pwd	:	[REDACTED]
Domain	:	egitlab	Admin_User	:	labadmin
Admin_Pwd	:	[REDACTED]	Confirm_Pwd	:	[REDACTED]
<input type="button" value="Update"/> <input type="button" value="Clear"/>					

Figure 4.10: 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 **jadm**. You can configure the eG agent with the credentials of both these users as shown by Figure 4.11.

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	: <span style="background-color: black; color: black;">*****</span>	Confirm_Pwd	: <span style="background-color: black; color: black;">*****</span>
Domain	: egitlab	Admin_User	: labadmin
Admin_Pwd	: <span style="background-color: black; color: black;">*****</span>	Confirm_Pwd	: <span style="background-color: black; color: black;">*****</span>
Domain	: egitlab	Admin_User	: jadmin
Admin_Pwd	: <span style="background-color: black; color: black;">*****</span>	Confirm_Pwd	: <span style="background-color: black; color: black;">*****</span>
<input type="button" value="Update"/> <input type="button" value="Clear"/>			

Figure 4.11: 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 4.11, 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., *jadmin* in our example (see Figure 4.11). 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 4.11.
8. To remove the details of a particular user alone, just click the  button in Figure 4.11.
9. To save the specification, just click on the **Update** button in Figure 4.11. 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 4.11).

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	: <input type="radio"/> Yes <input checked="" type="radio"/> No
DETAILED DIAGNOSIS	: <input checked="" type="radio"/> On <input type="radio"/> Off
APPLY TO OTHER COMPONENTS	: <input type="checkbox"/>

**Update**

Figure 4.12: The test configuration page displaying multiple domain names, user names, and passwords

#### 4.2.2 Disk Space - VM Test

This test monitors the space usage of every disk partition on a Windows virtual desktop.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every combination of *user\_on\_virtual\_desktop:disk\_partition*

#### Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is <i>NULL</i> .
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).

Parameter	Description
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.

### 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	
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.

### 4.2.3 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.

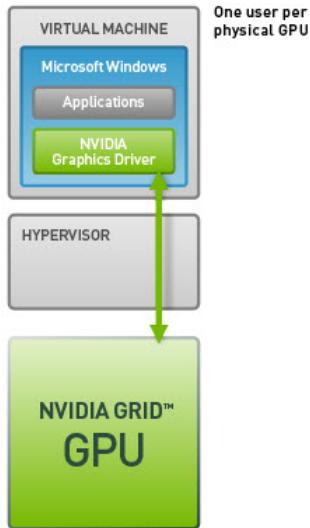


Figure 4.13: Dedicated GPU Technology

- **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.

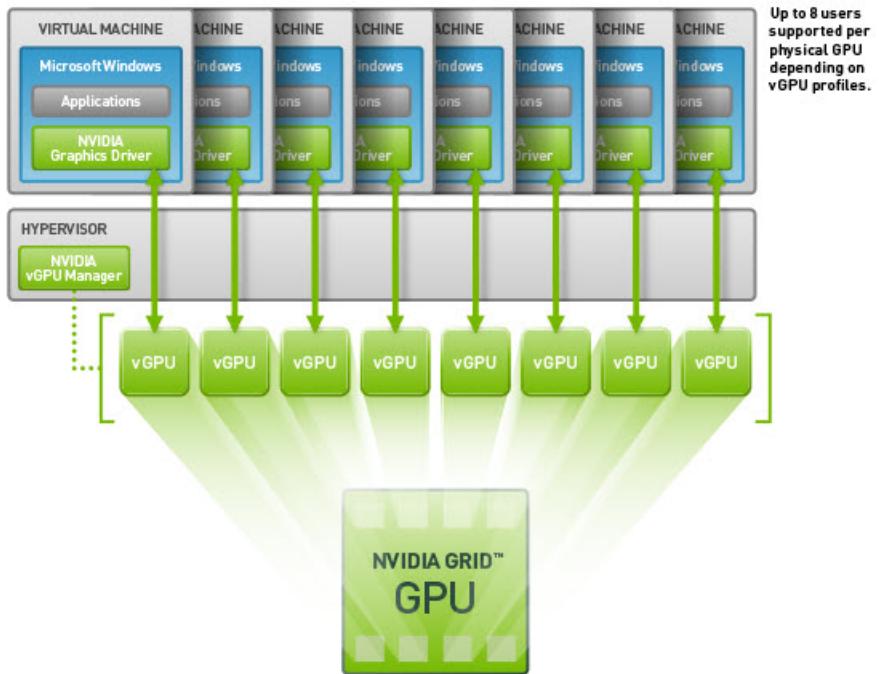


Figure 4.14: Shared vGPU Technology

In GPU-enabled VDI environments, if users to 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 virtual desktops? Is it because of excessive utilization of GPU memory and processing resources by a few virtual desktops? Or is it because the GPU clock frequencies are improperly set for one/more GPUs used by a virtual desktop? Accurate answers to these questions can help administrators determine whether/not:

- The 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 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 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 virtual desktops. In the process, the test also pinpoints those 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 virtual desktops where the **NVWMI** is installed. The steps for installing **NVWMI** and configuring the eG agent to use it have been detailed in [Configuring the eG Agent to Monitor NVIDIA Graphics Processing Units \(GPUs\)](#).

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results every vGPU assigned to each `user_on_virtualdesktop`.

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <code>username_on_virtualdesktopname</code> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.

Parameter	Description
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to <i>1:1</i> , it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
Detailed Diagnosis	<p>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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Cooler rate	Indicates the percentage of device cooler rate for this GPU of this virtual desktop.	Percentage	
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 virtual desktop.	Percentage	<p>A value close to 100% indicates that the GPU of the virtual desktop is busy processing graphic requests almost all the time.</p> <p>In a Shared vGPU environment a vGPU may be in use almost all the time, if the virtual desktop it is allocated to run graphic-intensive applications. A resource-hungry virtual desktop on a XenServer can impact the performance of other virtual desktops on the same server. If you find that only a single</p>

Measurement	Description	Measurement Unit	Interpretation
			<p>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 virtual desktop has no impact on the performance of other virtual desktops on that host.</p> <p>If all GPUs assigned to a virtual desktop are found to be busy most of the time, you may want to consider allocating more GPU resources to that virtual desktop.</p>
Power consumption	Indicates the current power usage of this GPU allocated to this virtual desktop.	Watts	<p>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.</p>
Core GPU temperature	Indicates the current temperature of this GPU allocated to this virtual desktop.	Celsius	<p>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 virtual desktops to identify that virtual desktop for which GPU temperature soared since the last reading.</p> <p>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.</p>

Measurement	Description	Measurement Unit	Interpretation
Total framebuffer memory	Indicates the total size of frame buffer memory of this GPU of this virtual desktop.	MiB	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 virtual desktop.	MiB	<p>Properties like the screen resolution, color level, and refresh speed of the frame buffer can impact graphics performance.</p> <p>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.</p> <p>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.</p> <p>For optimal graphics performance therefore, adequate frame buffer memory should be allocated to the virtual desktop.</p>
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.	MiB	
Frame buffer memory utilization	Indicates the percentage of total frame buffer memory that has been allocated to this VM/virtual desktop.	Percentage	<p>Ideally, the value of this measure should be low.</p> <p>A value close to 100% is indicative of excessive usage of frame buffer memory.</p> <p>Properties like the screen resolution, color level, and refresh speed of the</p>

Measurement	Description	Measurement Unit	Interpretation
			<p>frame buffer can impact graphics performance.</p> <p>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.</p> <p>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.</p> <p>For optimal graphics performance therefore, adequate frame buffer memory should be allocated to the VM/virtual desktop.</p>
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 period during which memory on this GPU was read/written on by this VM/virtual desktop.	Percentage	<p>A value close to 100% is a cause for concern as it indicates that the graphics memory on a GPU is almost always in use.</p> <p>In a Shared vGPU environment, memory may be consumed all the time if one/more 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 virtual desktops on</p>

Measurement	Description	Measurement Unit	Interpretation
			<p>that host.</p> <p>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.</p>
Total BAR1 memory	Indicates the total size of the BAR1 memory of this GPU allocated to this VM/virtual desktop.	MiB	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.	MiB	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.	MiB	
BAR1 memory utilization	Indicates the percentage of the allocated BAR1 memory that is currently being utilized by this VM/virtual desktop.	Percentage	<p>A value close to 100% is indicative of excessive usage of the BAR1 memory by a VM/virtual desktop.</p> <p>For best graphics performance, this value should be low. To ensure that, adequate BAR1 memory should be allocated to the VM.</p>
Power management	Indicates whether/not power management is enabled for this GPU of this VM/virtual desktop.		<p>Many NVIDIA graphics cards support multiple performance levels so that the server can save power when full graphics performance is not required.</p> <p>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</p>

Measurement	Description	Measurement Unit	Interpretation						
			<p>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.</p> <p>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.</p> <p>The numeric values that correspond to these measure values are discussed in the table below:</p> <table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Supported</td><td>1</td></tr> <tr> <td>Maximum</td><td>0</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure will report the <b>Measure Values</b> 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.</p>	Measure Value	Numeric Value	Supported	1	Maximum	0
Measure Value	Numeric Value								
Supported	1								
Maximum	0								
Power limit	Indicates the power limit	Watts	<b>This measure will report a value</b>						

Measurement	Description	Measurement Unit	Interpretation
	configured for this GPU of this VM/virtual desktop.		<p><b>only if the value of the ‘Power management’ measure is ‘Supported’.</b></p> <p>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.</p>
Default power limit	Indicates the default power management algorithm's power ceiling for this GPU.	Watts	<p><b>This measure will report a value only if the value of the ‘Power management’ measure is ‘Supported’.</b></p>
Enforced power limit:	Indicates the power management algorithm's power ceiling for this GPU of this VM/virtual desktop.	Watts	<p><b>This measure will report a value only if the value of the ‘Power management’ measure is ‘Supported’.</b></p> <p>The total board power draw is manipulated by the power management algorithm such that it stays under the value reported by this measure.</p>
Min power limit:	The minimum value that the power limit be set to for this GPU of this VM/virtual desktop.	Watts	<p><b>This measure will report a value only if the value of the ‘Power management’ measure is ‘Supported’.</b></p>
Max power limit:	The maximum value that the power limit for this GPU of this VM/virtual desktop can be set to.	Watts	<p><b>This measure will report a value only if the value of the ‘Power management’ measure is ‘Supported’.</b></p> <p>If the value of this measure is the same as that of the Power limit measure, then the GPU may behave strangely.</p>

Measurement	Description	Measurement Unit	Interpretation
Core clock:	Indicates current frequency of the graphics clock on this GPU of this VM/virtual desktop.	MHz	<p>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.</p> <p>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.</p> <p>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</p>

Measurement	Description	Measurement Unit	Interpretation
			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 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 of this VM.	Number	<p>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.</p> <p>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.</p>
Volatile single bit errors	Indicates the number of volatile single bit errors in this GPU of this VM/virtual desktop.	Number	<p>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.</p> <p>Ideally, the value of this measure should be 0.</p>
Volatile double bit errors	Indicates the total number of volatile double bit errors in this GPU of this VM/virtual desktop.	Number	<p>Volatile error counters track the number of errors detected since the last driver load. Double bit errors are detected but not corrected.</p> <p>Ideally, the value of this measure should be 0.</p>
Aggregate single bit errors	Indicates the total number of aggregate single bit errors in this GPU of this VM/virtual desktop.	Number	<p>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.</p>

Measurement	Description	Measurement Unit	Interpretation								
			<p>Ideally, the value of this measure should be 0.</p>								
Aggregate double bit errors	Indicates the total number of aggregate double bit errors in this GPU of this VM/virtual desktop.	Number	<p>Aggregate error counts persist indefinitely and thus act as a lifetime counter. Double bit errors are detected but not corrected.</p> <p>Ideally, the value of this measure should be 0.</p>								
Mode	Indicates the mode using which the GPU resources were delivered to this VMs.		<p>The values that this measure can take and their corresponding numeric values are as follows:</p> <table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Values</th></tr> </thead> <tbody> <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> </tbody> </table> <p><b>Note:</b></p> <p>By default, this test reports the <b>Measure Values</b> 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.</p>	Measure Value	Numeric Values	Pass through	0	Shared	1	Unavailable (GPU card is not allocated to any VM)	2
Measure Value	Numeric Values										
Pass through	0										
Shared	1										
Unavailable (GPU card is not allocated to any VM)	2										
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	<p><b>This measure will report metrics only VMs configured with a Tesla GPU card.</b></p> <p>A value close to 100% indicates that the physical GPU is busy processing graphic requests from this VM almost all the time.</p> <p>In a Shared vGPU environment a vGPU</p>								

Measurement	Description	Measurement Unit	Interpretation
			<p>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 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 virtual desktops on that host.</p> <p>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.</p>
Encoder utilization	Indicates the amount of physical GPU of this VM/virtual desktop that is utilized for encoding process.	Percent	<p><b>These measures will report metrics only VMs configured with a Tesla GPU card.</b></p> <p>A value close to 100 is a cause of concern. By closely analyzing these measures, administrators can easily be alerted to situations where graphics processing is a bottleneck.</p>
Decoder utilization	Indicates the amount of physical GPU of this VM/virtual desktop that is utilized for decoding process.	Percentage	

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

To determine whether/not the vGPUs assigned to a Windows virtual desktop are adequate for the graphic processing requirements of the that desktop, administrators must understand whether/not memory-intensive I/O operations are performed on that desktop and if so, how they impact vGPU

usage. To perform this check, administrators can periodically run the **GPU – VM** test on every Windows virtual desktop on the cloud.

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

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.

#### 4.2.3.1.1 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 4.15 and Figure 4.16 in the following procedure.

1. Download the full GPU driver from [www.nvidia.com/drivers](http://www.nvidia.com/drivers).
2. Double click the **setup.exe**.
3. Select **Custom (Advanced Options)** under “Installation options”.



Figure 4.15: Custom check box

4. Select the **NVIDIA WMI** check box under **Custom** installation options.



Figure 4.16: 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 Windows virtual desktop on which you want to run the **GPU Stats – VM** test.

#### 4.2.3.1.2 Installing NVWMI from the Standalone Installer

1. Go to [www.nvidia.com/drivers](http://www.nvidia.com/drivers).
2. Select the **NVIDIA Enterprise Manager** under **Beta, Older drivers and more**.

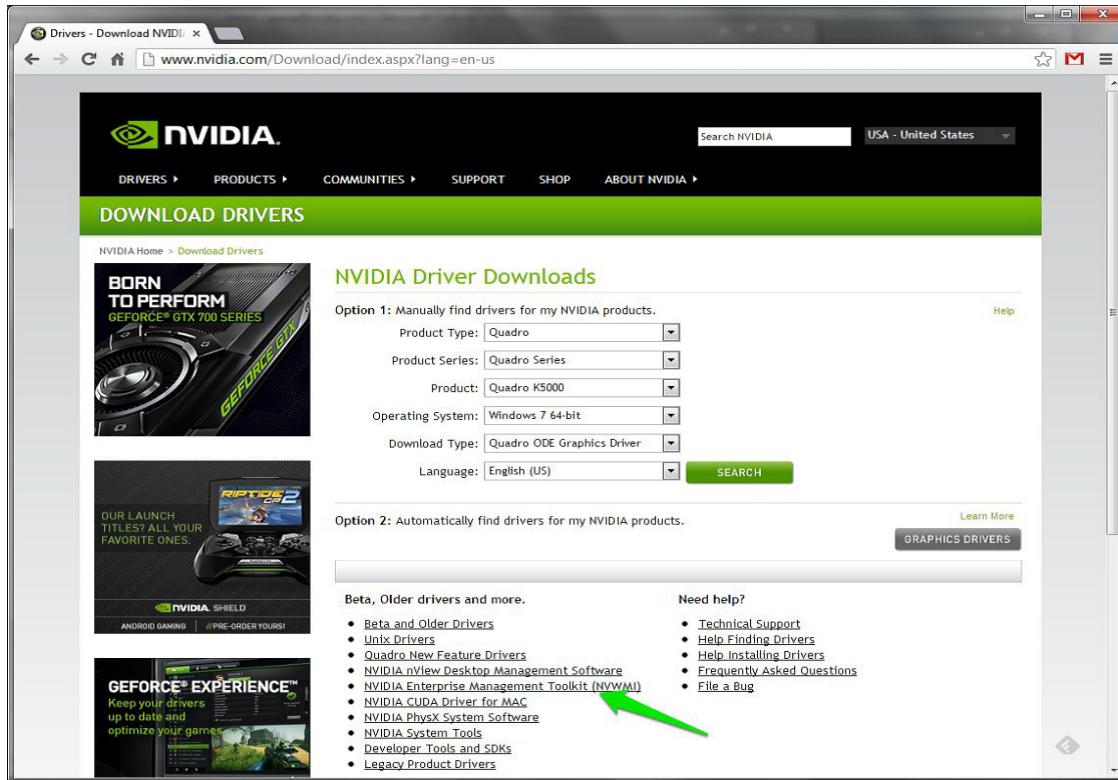


Figure 4.17: NVIDIA Enterprise Management Toolkit (NVWMI)

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

**DOWNLOAD**

<b>RELEASE HIGHLIGHTS</b>	<b>SUPPORTED PRODUCTS</b>	<b>ADDITIONAL INFORMATION</b>
---------------------------	---------------------------	-------------------------------

"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 NVVIEW 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 4.18: 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 Windows virtual desktop on which you want to run the **GPU Stats – VM** test.

#### 4.2.3.1.3 Installing NVWMI via the Command Line

1. Download the driver from the URL [www.nvidia.com/drivers](http://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
-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.

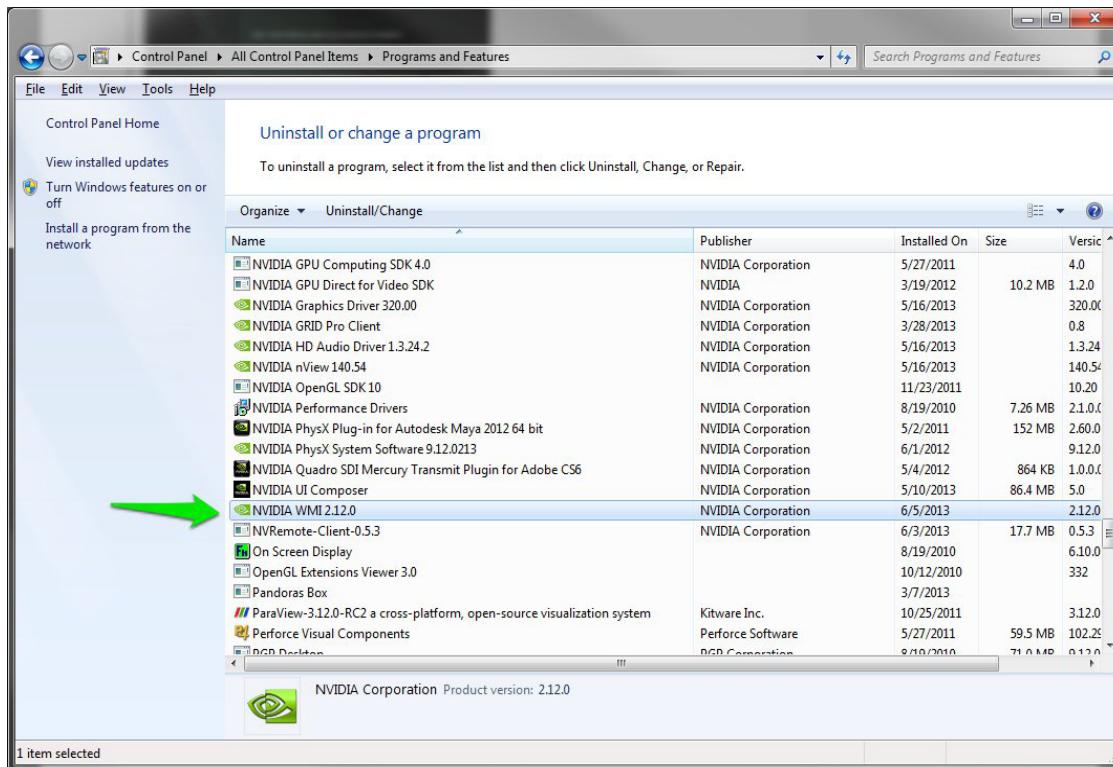


Figure 4.19: Checking the Control Panel for NVIDIA WMI

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

#### 4.2.4 System Details - VM Test

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

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every combination of *user\_on\_virtual\_desktop:processor*

### Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to <i>1:1</i> , it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"><li>• The eG manager license should allow the detailed diagnosis capability</li><li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li></ul>

### 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.
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 4.20). In the event of a CPU bottleneck, this information will enable users to identify the processes consuming a high percentage of CPU time. The users may then decide to stop such processes, so as to release the CPU resource for more important processing purposes.

Lists the top 10 CPU processes			
Time	PID	%CPU	ARGS
Jan 03, 2008 05:43:18	4	0.52	system
Jan 03, 2008 05:32:51	4	0.52	system
Jan 03, 2008 05:22:20	4	0.52	system
Jan 03, 2008 05:12:21	4	0.52	system
Jan 03, 2008 05:02:47	4	0.52	system
Jan 03, 2008 05:02:47	1768	0.52	xenservice
Jan 03, 2008 04:53:13	4	0.52	system

Figure 4.20: The top 10 CPU consuming processes

**Note:**

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

#### 4.2.5 Memory Usage - VM Test

This test reports statistics related to the usage of physical memory of each of the Windows virtual desktops.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every *user\_on\_virtualdesktop*

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to

Parameter	Description
	<p>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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### 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.
Physical memory utilized	Indicates the percent usage of physical memory by this VM.	Percent	<p>Ideally, the value of this measure should be low. While sporadic spikes in memory usage could be caused by one/more rogue processes on the VM, a consistent increase in this value could be a cause for some serious concern, as it indicates a gradual, but steady erosion of valuable memory resources. If this unhealthy trend is not repaired soon, it could severely hamper VM performance, causing anything from a slowdown to a complete system meltdown.</p> <p>You can use the detailed diagnosis of this measure to figure out which</p>

Measurement	Description	Measurement Unit	Interpretation
			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	<p>Not all of the Available physical memory is Free physical memory. Typically, Available physical memory is made up of the Standby List, Free List, and Zeroed List.</p> <p>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.</p> <p>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.</p> <p>All new allocations always come from the Zeroed List, which is memory pages that have been overwritten with zeros. This is a standard part of the OS' cross-process security, to prevent any process ever seeing data from another. If the Zeroed List is empty, Free List memory is zeroed and used or, if that is empty too, Standby List memory is freed, zeroed, and used. It is because all three can be used with so little effort that they are all counted as "available".</p> <p>A high value is typically desired for this</p>

Measurement	Description	Measurement Unit	Interpretation
			<p>measure.</p> <p><b>This measure will be available for Windows 2008 VMs only.</b></p>
Modified memory	Indicates the amount of memory that is allocated to the modified page list.	MB	<p>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.</p> <p>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.</p> <p><b>This measure will be available for Windows 2008 VMs only.</b></p>
Standby memory	Indicates the amount of memory assigned to the standby list.	MB	<p>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.</p> <p>Typically, Standby memory is the aggregate of Standby Cache Core Bytes,Standby Cache Normal Priority</p>

Measurement	Description	Measurement Unit	Interpretation
			<p>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.</p> <p><b>This measure will be available for Windows 2008 VMs only.</b></p>
Cached memory	This measure is an aggregate of Standby memory and Modified memory.	MB	<b>This measure will be available for Windows 2008 VMs only.</b>

#### 4.2.6 Handles Usage - VM Test

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

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for each *user\_on\_virtualdesktop*

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.

Parameter	Description
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
Handles Growth Limit	This defines the upper limit of the handles opened by any process. By default, this parameter is set to <b>100000</b> .
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to <b>1:1</b> , it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Handles used by processes	Indicates the number of handles opened by various processes	Number	Use the detailed diagnosis of this measure to determine the top-10 processes in terms of number of

Measurement	Description	Measurement Unit	Interpretation
	running in a target Windows virtual desktop in the last measurement period.		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 - <b>HANDLES GROWTH LIMIT</b> .	Number	<p>Using the detailed diagnosis of this measure, you can accurately isolate the process(es) that has opened more handles than the permitted limit.</p> <p>A high value of this measure indicates that too many processes are opening handles excessively. You might want to closely observe the handle usage of these processes over time to figure out whether the spike in usage is sporadic or consistent. A consistent increase in handle usage could indicate a handle leak.</p>

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
	cars	1097	420	468
	winlogon	564	420	492
	ImaSrv	559	540	3696
	Rtvscan	536	540	3936
	tomcat	485	540	6572
	services	482	492	540

Figure 4.21: 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	eGRSSvc	62410	412	11512

Figure 4.22: The detailed diagnosis of the Processes using handles above limit in VM measure

#### 4.2.7 PCoIP Session - VM Test

PCoIP - PC over IP - is a proprietary protocol for remote workstation and desktop resolution. Citrix 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 users who are currently connected to the Windows virtual desktops on the cloud 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.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every user who is connected to a Windows virtual desktop via PCoIP.

### Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.

### 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.

Measurement	Description	Measurement Unit	Interpretation
			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.
Imaging data sent rate	Indicates the bandwidth used when receiving imaging data from this user.	Kbit/Sec	
Imaging decoder capability rate	Indicates the current estimate of the decoder processing capability.	Kbit/Sec	
Incoming bandwidth	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	Kbit/Sec	

Measurement	Description	Measurement Unit	Interpretation
	channel.		
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	

#### 4.2.8 Blast Session - VM Test

The BLAST Session protocol allows files on a remote computer to be manipulated - i.e., to be deleted, renamed, or printed on the remote. Owing to the improved remote file management capabilities that Blast offers, many hypervisors like VMware support the Blast protocol too for user communication with virtual desktops. Because of this, there is a need to know which users have connected to virtual desktops via Blast and how the experience of each user is. This is why, the **Blast Session - VM** test auto-discovers the users who are connected to virtual desktops via the Blast protocol, and measures the bandwidth usage, frames processing ability, throughput, and time taken to establish each user's session. In the process, bandwidth-hungry, latent sessions can be accurately isolated.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every user who is connected to a Windows virtual desktop via the Blast protocol

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is <i>NULL</i> .

Parameter	Description
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Round trip time	Indicates the round trip latency between the virtual desktop and this user terminal.	Seconds	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 via Blast.
Bandwidth	Indicates the bandwidth usage of all sessions of this user.	Kbps	Compare the value of this measure across users to know which user is consuming the maximum bandwidth.
Frame rate	Indicates the rate at which frames are processed during this user's sessions.	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.
Throughput	Indicates the network throughput of this user's sessions.	Kbps	A poor network throughput could cause latencies in network communication.

#### 4.2.9 RemoteFX User Experience - VM Test

Microsoft RemoteFX™ enables the delivery of a full Windows user experience to a range of client devices including rich clients, thin clients, and ultrathin clients. RemoteFX delivers a rich user experience for Virtual Desktop Infrastructure (VDI) by providing a 3D virtual adapter, intelligent codecs, and the ability to redirect USB devices in virtual machines. RemoteFX is integrated with the RDP protocol, which enables shared encryption, authentication, management, and device support. RemoteFX also delivers a rich user experience for session-based desktops and RemoteApp programs to a broad range of client devices.

If a remote user's experience with a RemoteFX-enabled virtual desktop on the cloud is poor, then administrators should be able to quickly figure out what is causing the quality of the UX to suffer – is it poor frame quality? or severe packet loss? or bad picture output owing to a high compression ratio? or bottleneck in TCP/UDP connectivity? The **RemoteFX User Experience - VM** test helps answer this question. For each remote user connecting to a RemoteFX-enabled Windows virtual desktop, this test measures user experience and reports abnormalities (if any). This way, users who are experiencing a poor visual experience can be isolated and the reason for the same can be ascertained. In addition, the test points you to RemoteFX features that may have to be tweaked in order to improve overall performance.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every *user\_on\_virtual\_desktop*

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).

Report Powered OS **This flag is relevant only for those tests that are mapped to the Inside View of**

Parameter	Description
	<b>Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Average frames encoding time	Indicates the average time taken for encoding the frames of this user.	Secs	Compare the value of this measure across users to know for which user frames encoding took too long.
Frame quality	Indicates the quality of the output frame expressed as a percentage of the quality of the source frame for this user.	Percent	High frame rates produce a smooth representation of frames for the particular user, while low frame rates may cause rough or choppy representation of frames for the particular user. A high value is hence desired for this measure.  Compare the value of this measure across users to know which user received the poorest frame quality.
Frames skipped due to insufficient client resources	Indicates the rate at which frames were skipped for this user due to insufficient client resources.	Frames/Sec	A low value is desired for this measure. Compare the value of this measure across users to know which user is connecting from a client sized with inadequate resources.
Frames skipped due to insufficient network resources	Indicates the rate at which frames were skipped for this user due to insufficient network resources.	Frames/Sec	A low value is desired for this measure. Compare the value of this measure across users to know which user is connecting via a network that is sized with inadequate resources.
Frames skipped due	Indicates the rate at which	Frames/Sec	A low value is desired for this

Measurement	Description	Measurement Unit	Interpretation
to insufficient server resources	frames were skipped for this user due to insufficient server resources.		measure. Compare the value of this measure across users to know which user was unable to receive frames due to the lack of enough resources on the XenServer.
Graphics compression ratio	Indicates the ratio of the number of bytes encoded to the number of bytes input for this user.	Percent	The compression ratio typically affects the quality of the picture. Generally, the higher the compression ratio, the poorer the quality of the resulting picture. Ideally therefore, the value of this measure should be 0. You can compare the value of this measure across users to identify that user whose picture output was very poor owing to high compression.
Input frames	Indicates the number of source frames provided per second as input to the RemoteFx graphics for this user.	Frames/Sec	
Output Frames	Indicates the number of source frames sent per second to this user as output of RemoteFx graphics.	Frames/Sec	
Source frames	Indicates number of frames per second composed at the source for this user.	Frames/Sec	
Base TCP round trip time	Indicates the time between initiating a network request and receiving a response over TCP for this user.	Seconds	A high value for this measure could indicate a bottleneck in TCP connectivity between the user terminal and the server.
Base UDP round trip time	Indicates the time between initiating a network request and	Seconds	A high value for this measure could indicate a bottleneck in UDP connectivity between the user terminal

Measurement	Description	Measurement Unit	Interpretation
	receiving a response over UDP for this user.		and the server.
Current TCP bandwidth	Indicates the amount of data that is currently carried from one point to another over TCP for this user.	Kbps	A consistent rise in the value of this measure could indicate that TCP traffic to/from the user is consuming bandwidth excessively. Compare the value of this measure across users to identify that user who is performing bandwidth-intensive operations on the XenServer.
Current TCP round trip time	Indicates the average time between initiating a network request and receiving a response over TCP for this user.	Seconds	A high value could indicate a current problem with TCP connectivity between the user terminal and the server.
Current UDP bandwidth	Indicates the amount of data that is currently carried from one point to another over UDP for this user.	Kbps	A consistent rise in the value of this measure could indicate that UDP traffic to/from the user is consuming bandwidth excessively. Compare the value of this measure across users to identify that user who is performing bandwidth-intensive operations on the XenServer.
Current UDP round trip time	Indicates the average time between initiating a network request and receiving a response over UDP for this user.	Secs	A high value could indicate a current problem with UDP connectivity between the user terminal and the server.
Forward error correction	Indicates the percentage of forward error corrections performed for this user.	Percent	RemoteFX UDP transport uses Forward Error Correction (FEC) to recover from the lost data packets. In the cases where such packets can be recovered, the transport doesn't need to wait for the data to be retransmitted, which allows immediate delivery of data and prevents Head of Line Blocking. Preventing this stall results

Measurement	Description	Measurement Unit	Interpretation
			<p>in an overall improved responsiveness.</p> <p>A high value is hence desired for this measure.</p>
Loss	Indicates the percentage of packets lost when being transmitted to this user.	Percent	<p>A high value indicates that a large number of packets were lost without being retransmitted. By comparing the value of this measure across users, you can find that user who has suffered the maximum data loss. This could be owing to a bad network connection between the remote user terminal and the server.</p>
Retransmissions	Indicates the percentage of packets that have been retransmitted to this user.	Percent	<p>Retransmissions should only occur when it is certain that a packet to be retransmitted was actually lost. Redundant retransmissions can also occur because of lost acknowledgments, coarse feedback, and bad retransmissions.</p> <p>Retransmission rates over 5% can indicate degraded network performance on a LAN. The internet may vary between 5 and 15 percent depending upon traffic conditions. Any value above 25 percent indicates an excessive number of retransmissions that will significantly increase the time for the file transfer and annoy the user.</p>
TCP received rate	Indicates the rate at which the data is received over TCP for this user.	Kbps	<p>A high value is desired for these measures as it indicates high TCP throughput.</p>
TCP sent rate	Indicates the rate at which the data is sent over TCP for this user.	Kbps	
UDP received rate	Indicates the rate at which the data is received over UDP for this user.	Kbps	<p>A high value is desired for these measures as it indicates high UDP throughput.</p>

Measurement	Description	Measurement Unit	Interpretation
UDP sent rate	Indicates the rate at which the data is sent over UDP for this user.	Kbps	

#### 4.2.10 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 Windows virtual desktop on the cloud, and particularly monitors retransmissions.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for each *user\_on\_poweredonvirtualdesktop*

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and

Parameter	Description
	not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
Segments Sent Min	Specify the minimum threshold for the number of segments sent/transmitted over the network. The default value is 25; in this case, the test will compute/report the <i>Retransmit ratio from VM</i> measure only if more than 25 segments are sent over the network – i.e., if the value of the <i>Segments sent by VM</i> measure crosses the value 25. On the other hand, if the <i>Segments sent by VM</i> measure reports a value less than 25, then the test will not compute/report the <i>Retransmit ratio from VM</i> measure. This is done to ensure that no false alerts are generated by the eG Enterprise system for the <i>Retransmit ratio from VM</i> measure. You can change this minimum threshold to any value of your choice.

### 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

Measurement	Description	Measurement Unit	Interpretation
			buffer/queue overflows, or by lousy network links due to poor physical layer characteristics (e.g., low signal to noise ratio). By tracking the percentage of retransmissions at a guest, an administrator can quickly be alerted to problem situations in the network link(s) to the guest that may be impacting the service performance.

#### 4.2.11 Uptime - VM Test

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

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

This test included in the eG agent monitors the uptime of each Windows virtual desktop on the cloud.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every *user\_onPoweredOnVirtualDesktop*

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.

Parameter	Description
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
Report Manager Time	By default, this flag is set to <b>Yes</b> , indicating that, by default, the detailed diagnosis of this test, if enabled, will report the shutdown and reboot times of the virtual desktops in the manager's time zone. If this flag is set to <b>No</b> , then the shutdown and reboot times are shown in the time zone of the system where the agent is running (i.e., the system on which the remote agent is running).
Detailed Diagnosis	<p>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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

#### 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	Boolean	If this measure shows 1, it means that the guest was rebooted during the last measurement period. By checking the

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

#### 4.2.12 Tcp - VM Test

This test tracks various statistics pertaining to TCP connections to and from each Windows virtual desktop on the cloud. The details of the test are provided below:

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for each *user\_on\_poweredonvirtualdesktop*

## Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### 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 <i>ESTABLISHED</i> or <i>CLOSE_WAIT</i> states.
Connection drops on VM	Indicates the rate of established TCP connections dropped from the TCP listen queue.	Conns/Sec	This value should be 0 for most of the time. Any non-zero value implies that one or more applications on the guest are under overload.
Connection failures on VM	Indicates the rate of half open TCP connections dropped from the listen queue.	Conns/Sec	This value should be 0 for most of the time. A prolonged non-zero value can indicate either that the server is under SYN attack or that there is a problem with the network link to the server that is resulting in connections being dropped without completion.

#### 4.2.13 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 Windows virtual desktop (on the cloud), 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 Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every user who is currently logged into a Windows virtual desktop on the cloud.

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.

### 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	<p>Ideally, the value of this measure should be 100%.</p> <p>A low value for this measure denotes a high level of inactivity recently.</p>
Time since last activity	Indicates the time that has elapsed since this user performed an action on this desktop.	Mins	<p>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.</p>
Total idle time in session	Indicates the total time for which this user was idle during the session.	Mins	<p>If the value of this measure is the same as the value of the <i>Total time in session</i> measure for a user, it means that the user has been idle throughout the session.</p> <p>If the value of this measure is close to the value of the <i>Total time in session</i> measure for a user, it implies that the user has been idle for a long time.</p> <p>If the value of this measure is much lesser than the value of the <i>Total time in session</i> measure for a user, it means that the user has been active for most part of the session.</p>

#### 4.2.14 Windows Network Traffic - VM Test

This test monitors the incoming and outgoing traffic through each Windows virtual desktop on the cloud.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every *Windows\_virtual\_desktop:network\_interface* combination

### Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
Is Full Duplex	By default, this flag is set to <b>Yes</b> , 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 <b>No</b> . 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 <i>Incoming traffic</i> and <i>Outgoing traffic</i> 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	Mbps	An abnormally high rate of incoming traffic may require additional analysis.

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

**Note:**

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

#### 4.2.15 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 - VM test tracks critical metrics corresponding to the System PTEs and the pool areas of kernel memory of a Windows virtual desktop on the cloud.

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 *Cloud Desktops* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the < button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

**Target of the test :** A Windows desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every *user\_on\_poweredonvirtualdesktop*

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.

Parameter	Description
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.

### 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 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	

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

#### 4.2.16 Windows Security Center Status - VM Test

Windows Security Center (WSC) is a comprehensive reporting tool that helps administrators establish and maintain a protective security layer around Windows VMs to monitor the VM's health state. The Windows Security Center also monitors third party security products such as firewall, antivirus, antimalware and antispyware, installed on the VM. In order for the security products to be compliant with Windows and successfully report status to Action Center, these products should be registered with the security center. The security products communicate any subsequent status changes to the security center using private APIs. The security center, in turn, communicates these updates to Action Center, where they are finally displayed to the end user. With Windows Security Center, administrators can check whether any security product is installed and turned on, and if the definitions of the products are up to date and real-time protection is enabled. By continuously monitoring the Windows Security Center, administrators can instantly find out whether the security products are up-to-date or out dated, and the status of security products in real-time. This is what exactly the **Windows Security Center Status - VM** test does!

This test auto-discovers the security products installed on the Windows virtual desktops on the cloud, and for each security product reports the current definition status and the current protection status. Using these details, administrators are alerted to the systems on which the automatic updates are outdated and virus protection turned off. By closely monitoring the status, administrators can take necessary actions before the end users become vulnerable to virus threats or malicious attacks.

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 *Cloud Desktops* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the < button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every *security product:provider combination* on each *user\_on\_poweredonvirtualdesktop*

## Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to <i>1:1</i> , it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability.</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation								
Signature status	Indicates the current status of this security product.		<p>The values reported by this measure and its numeric equivalents are mentioned in the table below:</p> <table border="1" data-bbox="1019 559 1379 813"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Unknown</td><td>25</td></tr> <tr> <td>Up to date</td><td>15</td></tr> <tr> <td>Out of date</td><td>10</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports the <b>Measure Values</b> listed in the table above to indicate the current state of this security product. The graph of this measure however, represents the status of a server using the numeric equivalents only.</p> <p>Use the detailed diagnosis of this measure, to know about the name of Windows system on which the product is running, the file paths of product executables and the current status of the product.</p>	Measure Value	Numeric Value	Unknown	25	Up to date	15	Out of date	10
Measure Value	Numeric Value										
Unknown	25										
Up to date	15										
Out of date	10										
Real-time protection status	Indicates the real-time protection status of this security product.		<p>The values reported by this measure and its numeric equivalents are mentioned in the table below:</p>								

Measurement	Description	Measurement Unit	Interpretation	
			Measure Value	Numeric Value
Unknown	25			
Snoozed	20			
On	15			
Expired	10			
Off	0			

**Note:**

By default, this measure reports the **Measure Values** listed in the table above to indicate the current protection status of this security product. The graph of this measure however, represents the status of a server using the numeric equivalents only.

#### 4.2.17 Windows Services - VM Test

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

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every *user\_on\_poweredonvirtualdesktop*

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.

Parameter	Description
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to <i>1:1</i> , it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### 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	Number	The detailed diagnosis of this measure lists the services (with startup type as automatic) that are running.

Measurement	Description	Measurement Unit	Interpretation
	running in the 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.

#### 4.2.18 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 Windows virtual desktops on the cloud 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 *Cloud Desktops* 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 Windows virtual desktops being monitored should be managed by XenDesktop Broker.
- The Virtual Desktop Agent software should have been installed on the Windows virtual desktops.
- The **ICA Session** performance object should be enabled on the Windows virtual desktops.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A 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**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).

Parameter	Description
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.

### 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.
Audio bandwidth output	Indicates the bandwidth used while receiving sound/audio from this user.	Kbps	To minimize bandwidth consumption, you may want to consider disabling client audio mapping.
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 output	Indicates the bandwidth used when receiving data	Kbps	

Measurement	Description	Measurement Unit	Interpretation
	from this user's COM port.		
Drive bandwidth input	Indicates the bandwidth used when this user performs file operations on the mapped drive on the virtual desktop.	Kbps	Comparing the values of these measures 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.
Printer bandwidth output	Indicates the bandwidth used when the desktop responds to print jobs issued by this user.	Kbps	If bandwidth consumption is too high, you may want to consider disabling printing. Alternatively, you can avoid printing large documents over the ICA connection.
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 operations 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	Number	Compression reduces the size of the data that is transacted over the ICA

Measurement	Description	Measurement Unit	Interpretation
	ratio used from this user to the virtual desktop for a session.		channel.  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.
Session compression output	Indicates the compression ratio used from the virtual desktop to this user for a session.	Number	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.
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 for flash data bandwidth output	Indicates the bandwidth used from the virtual desktop to this user for flash data traffic	Kbps	

Measurement	Description	Measurement Unit	Interpretation
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 during this user session.	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.
Framehawk frame rate	Indicates the rate at which frames are processed by the Framehawk virtual channel, if it is enabled for	Frames/Sec	The Framehawk virtual channel optimizes the delivery of virtual desktops and applications to users on broadband wireless connections, when

Measurement	Description	Measurement Unit	Interpretation										
	this user session.		<p>high packet loss or congestion occurs.</p> <p><b>Note:</b></p> <p>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.</p>										
Framehawk network bandwidth	Indicates the bandwidth consumption of this user session when the Framehawk virtual delivery channel is used.	Kbps	<p>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.</p>										
Framehawk latency	Indicates the latency experienced by this user session when the Framehawk virtual delivery channel is used.	Secs	<p>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.</p>										
Framehawk network loss	Indicates the percentage of packet loss experienced by this user session when the Framehawk virtual delivery channel is used.	Percent	<p>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.</p>										
User's connection quality indicator	Indicates the connectivity of this user with the Citrix environment.		<p>The values that this measure can report and their corresponding numeric values are discussed in the table above:</p> <table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Poor connection</td><td>1</td></tr> <tr> <td>Weak connection</td><td>2</td></tr> <tr> <td>Strong connection</td><td>3</td></tr> <tr> <td>None</td><td>4</td></tr> </tbody> </table> <p><b>Note:</b></p>	Measure Value	Numeric Value	Poor connection	1	Weak connection	2	Strong connection	3	None	4
Measure Value	Numeric Value												
Poor connection	1												
Weak connection	2												
Strong connection	3												
None	4												

Measurement	Description	Measurement Unit	Interpretation
			<p>By default, this measure reports the <b>Measure Values</b> listed in the table above. In the graph of this measure however, the value of this measure is represented using their numeric equivalents only.</p> <p>A major alert will be generated when this test reports the "Poor connection" value for this measure. Likewise, a minor alert will be generated when the value of this measure is "Weak connection".</p>

#### 4.2.19 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 Windows virtual desktop on the cloud, 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 virtual desktop such personal vDisks are associated with, and what is consuming too much disk space – user profiles? Or user applications?

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 *Cloud Desktops* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the < button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for every *user\_on\_poweredonvirtualdesktop*

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to <i>1:1</i> , it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.

Parameter	Description
Detailed Diagnosis	<p>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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation								
Personal vDisk service status	Indicates whether Citrix Personal vDisk service is running or not on this VM.		<p>The values that this measure can report and their corresponding numeric values have been discussed in the table below:</p> <table border="1"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <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> </tbody> </table> <p><b>Note:</b></p> <p>By default, this test reports the <b>Measures Values</b> 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.</p>	Measure Value	Numeric Value	Stopped	0	Running	1	Not installed	2
Measure Value	Numeric Value										
Stopped	0										
Running	1										
Not installed	2										
Recompose status	Indicates the status of	Number	Use the detailed diagnosis of this								

Measurement	Description	Measurement Unit	Interpretation
	the initially provisioned disk or the updated image.		<p>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:</p> <ul style="list-style-type: none"> <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 partitioned by the service.</li> <li>• <b>Disk Format</b> – The personal vDisk is being formatted.</li> <li>• <b>Updating</b> – The initial provisioning or an image update is in progress.</li> <li>• <b>Error (Disk Discovery)</b> – An error state. An error occurred while discovering the personal vDisk.</li> <li>• <b>Error (Disk Init)</b> – An error state. An error occurred while partitioning or formatting the personal vDisk.</li> <li>• <b>Error (Sys Init)</b> – An error state. An error occurred while starting the Personal vDisk Service or configuring the personal vDisk.</li> <li>• <b>Error (Update)</b> – An error state. An error occurred during the initial provisioning or the last image update.</li> <li>• <b>Unknown</b> – An error state. An error occurred but the cause is unknown.</li> </ul>
Space used by user applications	Indicates the amount of space used by applications installed on	MB	Personal vDisks have two parts, which use different drive letters and are by default equally sized.

Measurement	Description	Measurement Unit	Interpretation
	the personal vDisk attached to this VM.		<p>One part comprises a Virtual Hard Disk file (a .vhdx file). This contains items such as applications installed in C:\Program Files. By default, this part uses drive V: but is hidden from users.</p> <p>These measures indicate how much space has been allocated to this .vhdx file and how much of the allocated space has been utilized by user applications contained in this file.</p>
Space allocated for user applications	Indicates the amount of space allocation for storing user applications on the personal vDisk attached to this VM.	MB	<p>A high value for the Space used by user applications and Space utilized by user applications measures is indicative of 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.</p>
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	<p>Personal vDisks have two parts, which use different drive letters and are by default equally sized.</p> <p>One part comprises C:\Users (in</p>

Measurement	Description	Measurement Unit	Interpretation
			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:.
Space allocated for user profiles	Indicates the amount of space allocated for storing user profiles on the personal vDisk attached to this VM.	MB	<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.</p> <p>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 user data on the personal vDisk .</p>
Free space	Indicates the amount of unused space on the personal vDisk attached to this VM.	MB	<p>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</p>

Measurement	Description	Measurement Unit	Interpretation
			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	<p>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.</p> <p>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 <b>Space utilized by user applications</b> and <b>Space utilized by user profiles</b> measures of that VM, to clearly understand what is occupying too much 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.</p>

#### 4.2.20 Disk Alignment - VM Test

In a SAN environment, the smallest hardware unit used by a SAN storage array to build a LUN out of multiple physical disks is called a chunk or a stripe. To optimize I/O, chunks are usually much larger than sectors. Thus, a SCSI I/O request that intends to read a sector in reality reads one chunk.

On top of this, in a Windows environment, NTFS is formatted in blocks ranging from 1MB to 8MB. The file system used by the guest operating system optimizes I/O by grouping sectors into so called clusters (allocation units).

Figure 4.23 shows these three layers at issue. There are the SAN blocks at the bottom, then the VMFS blocks in the middle, and then the NTFS blocks used by the Windows VM.



Figure 4.23: The SAN, VMFS, and NTFS blocks

If these three layers are not aligned, your SAN may be working harder than it needs to. For example, a call to read a single NTFS block may require the SAN to read three blocks as shown below:

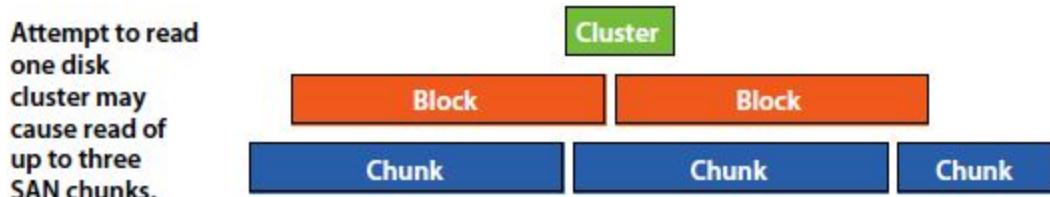


Figure 4.24: Unaligned partitions

An unaligned partition therefore, results in a track crossing and an additional I/O, incurring a penalty on latency and throughput. The additional I/O (especially if small) can impact system resources significantly on some host types.

What would hence be ideal is for the three layers in Figure 4.24 above to be aligned so that a single NTFS block requires only one SAN block to be read as illustrated below:

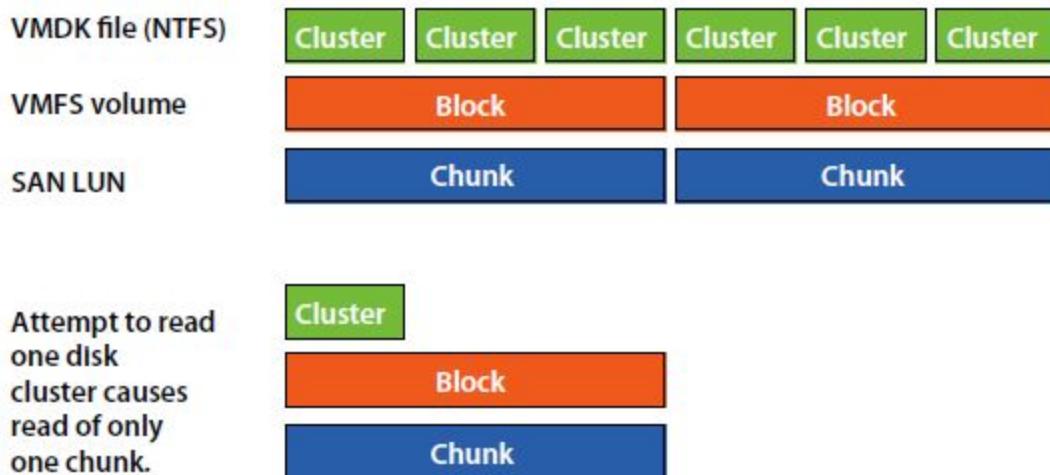


Figure 4.25: Aligned partitions

An aligned partition ensures that the single I/O is serviced by a single device, eliminating the additional I/O and resulting in overall performance improvement.

Therefore, whenever users to Windows virtual desktops complaint that the virtual desktop is running slower than usual, you may want to check the disk alignment to determine whether the slowdown can be attributed to one/more unaligned disk partitions. This test enables you to perform such a check.

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 *Cloud Desktops* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the < button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** One set of results for each disk partition on every Windows virtual desktop on the cloud

#### Configurable parameters for the test

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

Parameter	Description
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to <i>1:1</i> , it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.  The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: <ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Disk partition alignment status	Indicates whether this disk partition is aligned or not.		If the partition is unaligned, this test reports the value <i>Partition is not aligned</i> . For an aligned partition, this

Measurement	Description	Measurement Unit	Interpretation						
			<p>test reports the value <i>Partition is aligned</i>.</p> <p>The numeric values that correspond to the above-mentioned measure values are described in the table below:</p> <table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Partition is aligned</td><td>100</td></tr> <tr> <td>Partition is not aligned</td><td>0</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure reports one of the <b>Measure Values</b> listed in the table above. The graph of this measure however will represent the disk alignment status using the numeric equivalents - 100 or 0.</p> <p>If a partition is found to be misaligned, you can use the detailed diagnosis of this test to figure out the caption, device ID, logical partition name, and block size of the faulty partition.</p>	Measure Value	Numeric Value	Partition is aligned	100	Partition is not aligned	0
Measure Value	Numeric Value								
Partition is aligned	100								
Partition is not aligned	0								

If a partition is found to be misaligned, you can use the detailed diagnosis of this test to figure out the caption, device ID, logical partition name, and block size of the faulty partition.

Details of physical disk partition				
Time	Caption	Device id	Logical partition name	Block size
May 31, 2011 17:38:37	VMware Virtual disk SCSI Disk Device	\\.\PHYSICALDRIVE0	C:	512

Figure 4.26: The detailed diagnosis of the Disk partition alignment status measure

## 4.2.21 Page File - VM Test

When the load imposed by applications and services running on a server nears the amount of installed RAM, additional storage is necessary. The page file serves as the temporary store on disk for memory that cannot be accommodated in the physical RAM. Since it is frequently accessed for storing and retrieving data that is needed for virtual memory access by application, the location and sizing of the page files can have a critical impact on server's performance. Ideally, the server operating system and the page file should be available on different drives for optimal performance. Splitting the page file across different drives can improve performance further.

A rule of thumb in sizing the page file is to set the maximum size of the page file to 1.5 times the available RAM. While this works well for systems with smaller physical memory, for other systems, the optimal page file size has to be determined based on experience using the system and studying the typical workload.

This test tracks the usage of each of the page files on a Windows virtual desktop.

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 *Cloud Desktops* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the < button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent executing the test :** A remote agent

**Output of the test :** one set of results for each page file on each Windows virtual desktop

**Configurable parameters for the test**

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).

Parameter	Description
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
Report Total	Set this flag to <b>Yes</b> if you want the test to report total page file usage - i.e., the aggregate usage across multiple page files. In this case therefore, a <b>Total</b> descriptor will newly appear for this test in the eG monitoring console.
Report Total Only	If both the <b>REPORTTOTAL</b> and <b>REPORTTOTALONLY</b> flags are set to <b>Yes</b> , then the test will report only the aggregate usage across multiple page files - in other words, the test will report values for the <b>Total</b> descriptor only. Likewise, if the <b>REPORTTOTAL</b> flag is set to <b>No</b> , and the <b>REPORTTOTALONLY</b> flag is set to <b>Yes</b> , then again, the test will report current usage for the <b>Total</b> descriptor only. However, if both the <b>REPORTTOTAL</b> and <b>REPORTTOTALONLY</b> flags are set to <b>No</b> , then the test will report individual usages only. Also, if the <b>REPORTTOTAL</b> flag is set to <b>Yes</b> and the <b>REPORTTOTALONLY</b> flag is set to <b>No</b> , then both the individual and Total usages will be reported.

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current usage	Indicates the current usage of a page file.	Percent	This metric should be less than 90%. If the page file does not have additional space, additional users/processes cannot be supported and system performance will suffer. To improve performance, consider resizing the page file. Microsoft Windows allows a minimum and maximum size of the page file to be specified. If the system has sufficient disk space, consider setting the page file to start out at the maximum size (by using

Measurement	Description	Measurement Unit	Interpretation
			the same value for the minimum and maximum sizes), so that system resources are not spent growing the page file size when there is a virtual memory shortage.

#### 4.2.22 Windows Update Details - VM Test

Microsoft regularly releases various Windows updates to enhance and protect the Windows operating system. These updates are also applicable for the virtual desktops that host Windows operating system. The Windows updates fix newly discovered security holes and bugs, add malware definitions to Windows Defender and Security Essentials utilities, strengthen Office security and add new features/enhancements to the Windows operating system. By installing these updates regularly, you can keep the operating system highly secure, reliable and stable, and can maintain the performance of the operating system at peak. If the operating system is not updated regularly, the critical bugs and security errors may increase vulnerabilities. These vulnerabilities can be exploited by the malware or hackers, thus exposing the operating system to malicious attacks and degrading the operating system's performance. To avoid such eventualities, you should regularly check whether the Windows operating system is up-to-date or not. This check can be easily done using the **Windows Update Details - VM** test.

This test continuously monitors the virtual desktops and reports the current status of the Windows updates on the Windows operating system hosted by the virtual desktops. Besides, this test indicates whether any update is pending for the operating system and whether the Windows system is rebooted or not. In the process, this test also reports the total number of updates to be installed for the operating system and the number of Windows updates of different types at regular intervals.

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 *Cloud Desktops* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the < button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

**Target of the test :** A Windows virtual desktop on the cloud

**Agent deploying the test :** A remote agent

**Outputs of the test :** one set of results for each Windows virtual desktop.

## Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The nick name of the Cloud Desktops component for which this test is to be configured.
Port	Refers to the port at which the specified host listens to. By default, this is NULL.
Inside View Using	To obtain the 'inside view' of performance of the cloud-hosted Windows desktops - i.e., to measure the internal performance of the Windows virtual desktops - this test uses a light-weight eG VM Agent software deployed on each of the desktops. Accordingly, this parameter is by default set to eG VM Agent (Windows).
Report Powered OS	<b>This flag is relevant only for those tests that are mapped to the Inside View of Desktops layer.</b> If this flag is set to <b>Yes</b> (which is the default setting), then the 'inside view' tests will report measures for even those Windows virtual desktops that do not have any users logged in currently. Such desktops will be identified by their name and not by the <i>username_on_virtualdesktopname</i> . On the other hand, if this flag is set to <b>No</b> , then this test will not report measures for those Windows virtual desktops to which no users are logged in currently.
Is Cloud VMs	Since this test runs for a 'Cloud Desktops' component, this flag is set to <b>Yes</b> by default.
DD For Total Updates	In large VDI environments where hundreds of Windows virtual desktops have been provisioned, the frequent collection of detailed diagnosis information related to the update details of the virtual desktops may increase the processing overheads of the eG agent, and may even choke the eG database. To avoid this, by default, the DD For Total Updates flag is set to <b>No</b> indicating that this test will not report the detailed diagnostics for the <i>Total Updates Available</i> measure. However, you can set this flag to <b>Yes</b> if you want to collect the detailed diagnostics of the <i>Total Updates Available</i> measure.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For instance, if you set to <i>1:1</i> , it means that detailed measures will be generated every time this test runs, and also every time the test detects a problem.
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 <b>On</b> option. To disable the capability, click on the <b>Off</b> option.
	The option to selectively enable/disable the detailed diagnosis capability will be

Parameter	Description
available only if the following conditions are fulfilled:	
<ul style="list-style-type: none"> <li>• The eG manager license should allow the detailed diagnosis capability.</li> <li>• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</li> </ul>	

### Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation						
Are pending updates available?	Indicates whether/not the updates are pending.		<p>The values that this measure can report and the numeric values they indicate have been listed in the table below:</p> <table border="1"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>No</td> <td>0</td> </tr> <tr> <td>Yes</td> <td>1</td> </tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure can report the <b>Measure Values</b> mentioned above while indicating whether/not the updates are available. However, the graph of this measure is indicated using the numeric equivalents.</p>	Measure Value	Numeric Value	No	0	Yes	1
Measure Value	Numeric Value								
No	0								
Yes	1								
Is a system reboot pending?	Indicates whether the Windows system is rebooted or not.		<p>The values that this measure can report and the numeric values they indicate have been listed in the table below:</p> <table border="1"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>No</td> <td>0</td> </tr> <tr> <td>Yes</td> <td>1</td> </tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure can report the</p>	Measure Value	Numeric Value	No	0	Yes	1
Measure Value	Numeric Value								
No	0								
Yes	1								

Measurement	Description	Measurement Unit	Interpretation																		
			<p><b>Measure Values</b> mentioned above while indicating whether the system is rebooted or not. However, the graph of this measure is indicated using the numeric equivalents.</p>																		
Windows update service status	Indicates the current status of the Windows update service.		<p>The values that this measure can report and the numeric values they indicate have been listed in the table below:</p> <table border="1"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>Unknown</td><td>0</td></tr> <tr> <td>Running</td><td>1</td></tr> <tr> <td>Start pending</td><td>2</td></tr> <tr> <td>Continue pending</td><td>3</td></tr> <tr> <td>Pause pending</td><td>4</td></tr> <tr> <td>Stop pending</td><td>5</td></tr> <tr> <td>Paused</td><td>6</td></tr> <tr> <td>Stopped</td><td>7</td></tr> </tbody> </table> <p><b>Note:</b></p> <p>By default, this measure can report the <b>Measure Values</b> mentioned above while indicating the current status of Windows update service. However, the graph of this measure is indicated using the numeric equivalents.</p>	Measure Value	Numeric Value	Unknown	0	Running	1	Start pending	2	Continue pending	3	Pause pending	4	Stop pending	5	Paused	6	Stopped	7
Measure Value	Numeric Value																				
Unknown	0																				
Running	1																				
Start pending	2																				
Continue pending	3																				
Pause pending	4																				
Stop pending	5																				
Paused	6																				
Stopped	7																				
Total updates available	Indicates the total number of Windows updates available for the operating system.	Number	The detailed diagnosis of this measure, if enabled, lists the Windows updates available for the system and the categories of the available updates.																		
Critical updates available	Indicates the number of critical updates available for the operating system.	Number	A critical update is a widely and frequently released update that deals with the specific, non-security related, critical bugs. If these bugs are not fixed																		

Measurement	Description	Measurement Unit	Interpretation
			quickly, they can cause serious performance degradation, interoperability malfunction or disturb application compatibility.
Important updates available	Indicates the number of important updates available for operating system.	Number	The important updates help fixing the vulnerabilities using which malware/hackers can exploit the system resources or steal data. This in turn may leave the confidentiality and integrity of the system defenseless and make the user data unavailable.
Moderate updates available	Indicates the number of moderate security updates available for the operating system.	Number	The moderate updates fix a vulnerability whose exploitation can be mitigated to a significant degree by default configuration, auditing, or difficulty of exploitation.
Low updates available	Indicates the number of low security updates available for the operating system.	Number	These updates fix the vulnerability whose exploitation is extremely difficult.
Optional updates available	Indicates the number of optional updates available for operating system.	Number	An optional update includes Feature Pack and standard Updates, and does not have a severity rating.