Monitoring VDI-In-a-Box

eG Innovations Product Documentation

www.eginnovations.com

Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

Citrix VDI-in-a-Box is a simple, affordable and centrally managed desktop virtualization solution that can be used by businesses of all sizes. Users can login to the virtual desktops through a thin client and access the same applications as they did on physical desktops. Enterprises look at Citrix VDI-in-a-Box to simplify management of desktops and lower the total cost of ownership of desktops.

For a VDI-in-a-Box deployment to be successful, users should see the same performance when accessing their virtual desktops as they did when they accessed applications on a physical desktop. Hence, performance monitoring of VDI-in-a-Box is important, in order to ensure that virtual desktops deliver the same or better performance than physical desktops.

Performance monitoring for VDI-in-a-Box is challenging because even though VDI-in-a-Box offers a simple usage model, its underlying infrastructure is heterogeneous, multi-tier and interdependent. A typical VDI-in-a-Box deployment includes one or more physical servers hosting a virtualization platform (Citrix XenServer, VMware vSphere or Microsoft Hyper-V), the VDI-in-a-Box appliance(s) which includes the connection broker, load balancer, user manager, and desktop provisioning server, local storage on the hypervisors and the virtual desktops (see Figure 1). Multiple VDI-in-a-Box appliances can be linked in a grid for scalability and redundancy.

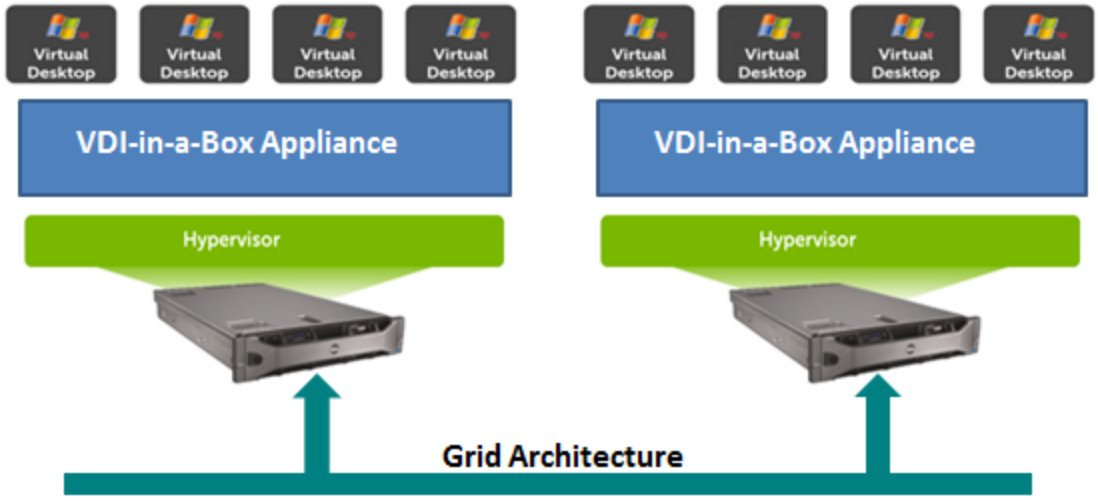

Figure 1.1: A typical VDI-in-a-Box infrastructure

To access their virtual desktops, users connect from their client devices to the VDI-in-a-box broker. The broker authenticates the user using Active Directory and then provisions a virtual desktop.

When a performance problem happens, IT managers need to be able to quickly determine where is the slowdown – is it due to the network? Or the hypervisor? Or the VDI-in-a-Box broker? Or the local storage? Or could it be due to one of the applications running in the virtual desktop?

Often, virtual desktop administrators use monitoring tools available with the hypervisor for monitoring virtual desktop infrastructures. Here are three main reasons why a virtual platform monitoring tool like VMware vCenter or Citrix XenCenter is not sufficient for monitoring VDI-in-a-Box deployments:

- **Monitor users, not just VMs**: The workload of a VM (virtual desktop) depends upon which user is accessing a virtual desktop and what activity he or she is performing on the virtual desktop. Therefore, VDI-in-a-Box monitoring must be based on which user is logged into a VM and not just based on the VM name itself. Virtualization platform monitoring tools do not have this capability.
- **Monitor activity inside the desktops**: Virtualization platform monitoring tools provide the "outside" view of a virtual desktop – i.e., what CPU or memory resources is a desktop consuming. This outside view of a virtual desktop does not provide details into why that desktop is using excessive resources – is it because of one malfunctioning application, or is it due to some activity the user is performing, etc. A VDI-in-a-Box monitoring solution must be able to monitor activity inside each and every virtual desktop. Furthermore, since tens to hundreds of desktops could be hosted on a physical server, the VDI-in-a-Box monitoring solution should not require agents on each and every virtual desktop to provide this "inside view" of virtual desktops.
- **Monitor the VDI-in-a-Box service infrastructure end-to-end**: A VMware vCenter or a Citrix XenCenter only tracks the health of the virtual platform. A VDI-in-a-Box monitoring solution must be able to monitor every layer and every tier of the virtual desktop infrastructure. It must provide visibility into the hypervisor, the physical server hardware, the Active Directory servers, the VDI-in-a-Box appliance, the local storage, the interconnecting networks and the virtual desktops themselves.

The eG VDI-in-a-Box Monitor is a turnkey monitoring, diagnosis and reporting solution for VDI-in-a-Box infrastructures. It offers dedicated monitoring models for each of the hypervisors that VDI-in-a-Box supports; these models are, namely – VDI in a Box / VMware, VDI in a Box / XenServer, VDI in a Box / Hyper-V. These models provide detailed insights into all the critical tiers of the VDI-in-Box infrastructure – i.e., the physical server, the hypervisor, the virtual desktops, and the VDI-in-a-Box appliance - thus enabling administrators to quickly spot current/probable slowdowns in the virtual desktop service and accurately pinpoint the reasons for the same.

1.1 What the eG VDI-in-a-Box Monitor Reveals?

Once the pre-requisites are fulfilled, the eG agent will be able to pull out a variety of performance information, with the help of which the following can be ascertained:

Monitored Entity	Details of Monitoring
VDI-in-a-Box Manager	• Is the VDI-in-a-Box appliance powered on and available? • Are the key VDI-in-a-Box processes (tomcat, license manager, etc.) running? • Is the VDI-in-a-Box manager accessible over HTTP/S? • If so, how quickly is the VDI-in-a-Box manager responding to HTTP/S connection requests? • Are the key TCP ports used for communication with other VDI-in-a-Box managers in a grid open? • How many VDI-in-a-Box managers are in the grid and what is their state? • Is any of the servers in the grid seeing a performance bottleneck?
VDI-in-a-Box Licenses	• How many virtual desktop licenses are installed and how many are in use? • Are sufficient licenses available to provision additional virtual desktops?
Virtualization Platform	• Is the hardware working well? • Is the hypervisor resource bottlenecked – for CPU or memory? • What is the space utilization of the local data stores? • Is any of the storage LUNs seeing excessive I/O operations? Which one? • Are the network interfaces working well? • How many VMs are powered on simultaneously? • Is any VM taking excessive resources?
Virtual Desktops	• What is the status of the virtual desktops configured on the VDI-in-a-Box broker? • How many virtual desktops are currently in use on a physical server?

Monitored Entity	Details of Monitoring
	• Is any virtual desktop using excessive amount of CPU or memory or causing a lot of IOPS? • If so, what application processes within the virtual desktop are responsible for the resource consumption? • Is there excessive queuing for disk access on any of the Virtual Desktops? • What are peak usage times of the virtual desktops? • Is there sufficient disk space in each of the disk partitions of the virtual desktop?
Users and Sessions	• Who are the users logged on to the Virtual Desktops? • What is the current state of a user's desktop session? • When did a user login and how long did he/she remain logged in for? • Who are the most resource intensive users?

The chapters that follow will discuss each of these monitoring models in detail.

Chapter 2: How to Monitor VDI-in-a-Box Using eG Enterprise?

eG VDI-in-a-Box Monitor model monitors the VDI-in-a-Box in an agentless manner. For this purpose, you require a single eG agent to be deployed on a remote Windows/Linux host in your environment. This remote agent should be configured to:

- Connect to the vdimanager via SSH and pull out metrics related to the overall health and availability of the vdimanager and the usage of the virtual desktops it manages;
- Connect to the physical server on which the vdimanager operates, and report on the health and resource usage of the host, the hypervisor, and the virtual desktops using a patented 'In-N-Out' monitoring algorithm.

To enable the eG agent to pullout metrics, a set of pre-requisites should be fulfilled. These requirements are discussed in the section below.

2.1 Pre-requisites for monitoring VDI-in-a-Box?

To enable the eG remote agent to pull out all these performance statistics, you first need to make sure that the following pre-requisites are in place:

For monitoring the vdimanager:

- The eG agent should be able to connect to the vdimanager via SSH and pull out the metrics. For this purpose, the default SSH port, 22, should be opened on the vdimanager. If your environment has been configured with a different SSH port, then make sure that that port is open.
- The eG agent should be able to login to the vdimanager appliance for monitoring and metrics collection. To enable this, you need to configure all tests that the eG agent executes on the vdimanager with the credentials of a user with login rights. By default, the appliance supports a root user named root and a user named **kvm**. If you prefer not to expose the credentials of the root user, then you can configure the tests with the credentials of the other user, **kvm**, for this purpose. By default, the user root takes the password root, and the user **kvm** takes the password **kaviza123**.

For 'In-N-Out' monitoring of the VMware vSphere server and its VMs:

Make sure that the pre-requisites detailed in Section 2.1 of the Monitoring VMware Infrastructures document are fulfilled.

For 'In-N-Out' monitoring of the Citrix XenServer and its VMs:

Make sure that the pre-requisites pertaining to 'agentless' monitoring of XenServers (detailed in Section 1.3.2 of the Monitoring XenServers document) are fulfilled.

For 'In-N-Out' monitoring of the Microsoft Hyper-V and its VMs:

Make sure that the pre-requisites pertaining to 'agentless' monitoring of the Hyper-V server (detailed in Section 1.3 of the Monitoring Microsoft Hyper-V document) are fulfilled.

2.2 Managing the VDI-in-a-Box on VMware vSphere

The eG Enterprise cannot automatically discover the *VDI in a Box / VMware* so that you need to manually add the component for monitoring. Remember that the eG Enterprise automatically manages the components that are added manually. To manage a *VDI in a Box / VMware* component, do the following:

1. Log into the eG administrative interface.
2. Follow the Components -> Add/Modify menu sequence in the **Infrastructure** tile of the **Admin** menu.
3. In the **COMPONENT** page that appears next, select *VDI in a Box / VMware* as the **Component type**. Then, click the **Add New Component** button. This will invoke 2.2.

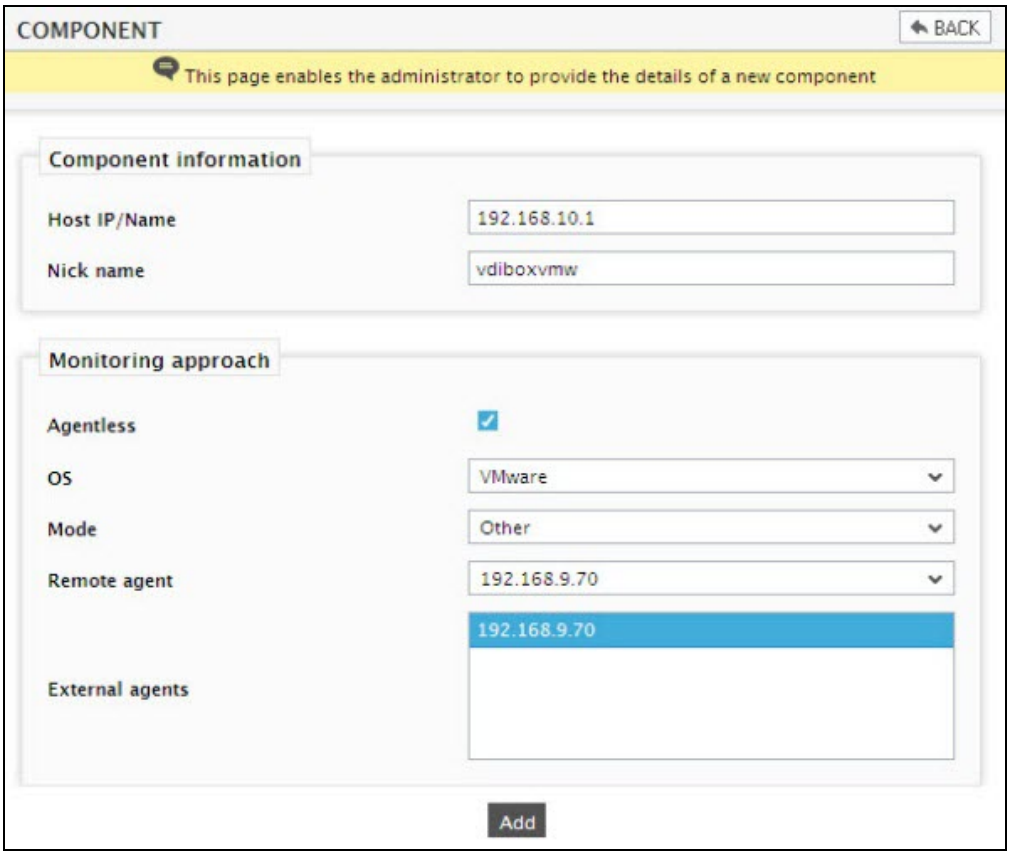

Figure 2.1: Adding a new VDI in a Box / VMware component

4. Specify the **Host IP/Name** and the **Nick name** of the VDI in a Box / VMware in 2.2. Then click on the Add button to register the changes.
5. Next, when you try to signout of the eG administrative interface, a list of unconfigured tests will appear as shown in Figure 2.2.

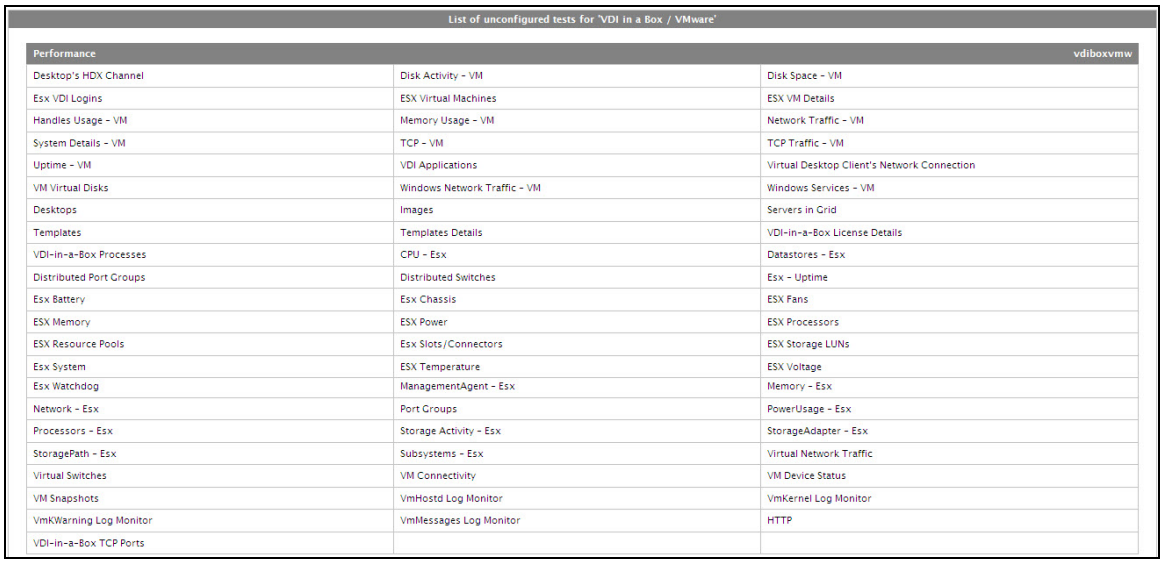

List of unconfigured tests for 'VDI in a Box / VMware'

Performance		vdiboxvmw
Desktop's HDX Channel	Disk Activity - VM	Disk Space - VM
Esx VDI Logins	ESX Virtual Machines	ESX VM Details
Handles Usage - VM	Memory Usage - VM	Network Traffic - VM
System Details - VM	TCP - VM	TCP Traffic - VM
Uptime - VM	VDI Applications	Virtual Desktop Client's Network Connection
VM Virtual Disks	Windows Network Traffic - VM	Windows Services - VM
Desktops	Images	Servers in Grid
Templates	Templates Details	VDI-in-a-Box License Details
VDI-in-a-Box Processes	CPU - Esx	Datastores - Esx
Distributed Port Groups	Distributed Switches	Esx - Uptime
Esx Battery	Esx Chassis	ESX Fans
ESX Memory	ESX Power	ESX Processors
ESX Resource Pools	Esx Slots/Connectors	ESX Storage LUNs
Esx System	ESX Temperature	ESX Voltage
Esx Watchdog	ManagementAgent - Esx	Memory - Esx
Network - Esx	Port Groups	PowerUsage - Esx
Processors - Esx	Storage Activity - Esx	StorageAdapter - Esx
StoragePath - Esx	Subsystems - Esx	Virtual Network Traffic
Virtual Switches	VM Connectivity	VM Device Status
VM Snapshots	VmHostd Log Monitor	VmKernel Log Monitor
VmKWarning Log Monitor	VmMessages Log Monitor	HTTP
VDI-in-a-Box TCP Ports		

Figure 2.2: Viewing the list of tests to be configured for the VDI in a Box / VMware

3. Click on any test in the list of unconfigured tests.
4. Now when you try to signout of the admin interface, you will be prompted to configure a list of tests. Try to configure the Disk Activity – VM test as shown in Figure 1.4.
5. To know more on how to configure the **Disk Activity – VM** test, Refer to the *Monitoring VMware Infrastructures* document.
6. Finally, you will be prompted to configure the HTTP test when you try to signout of the eG administrative interface. Refer to the *Monitoring VDI in a Box* document to know more about how to configure the HTTP test.
7. With that, test configuration is complete and the VDI in a BOX / VMware is finally ready to be monitored.

Note:

You can configure and monitor the **VDI in a Box / XenServer** and **VDI in a Box / Hyper-V** models using the same procedure as discussed above for the **VDI in a Box / VMware**.

Chapter 3: Monitoring VDI-in-a-Box on VMware vSphere

To monitor VMware vSphere and the vdimanager operating on it, use the VDI in a Box / VMware monitoring model.

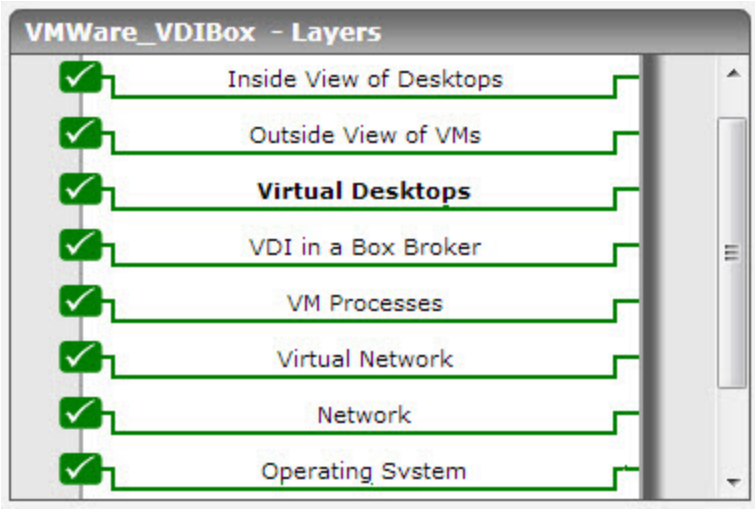

Figure 3.1: Layer model of VDI in a Box/VMware

Each layer of this model is mapped to tests that report a variety of statistics related to the VMware vSphere server and the VDI-in-a-Box appliance running on it. While the five layers at the bottom of Figure 2.1 focus on the health of the hardware, operating system, and the network of the VMware host, the top two layers perform 'In-N-Out' monitoring of the virtual desktops operating on the target VMware host. To ascertain the status of the vdimanager and its operations on the other hand, you need to use the Virtual Desktops and VDI in a Box Broker layers.

The sections that follow will discuss the Virtual Desktops and vdi In a Box Broker layers alone, as all the other layers have been dealt with elaborately in the Monitoring VMware Infrastructures document.

3.1 The VDI in a Box Manager Layer

The tests mapped to this layer, reveal the following:

- The availability and responsiveness of the VDI-in-a-Box console;
- The usage of the VDI-in-a-Box licenses;
- The health of the servers in a grid;

- Whether processes critical to the functioning of the broker are available or not, and if available, how is their resource usage;
- Whether important TCP ports are available on the appliance or not, and if so, how responsive they are to TCP requests

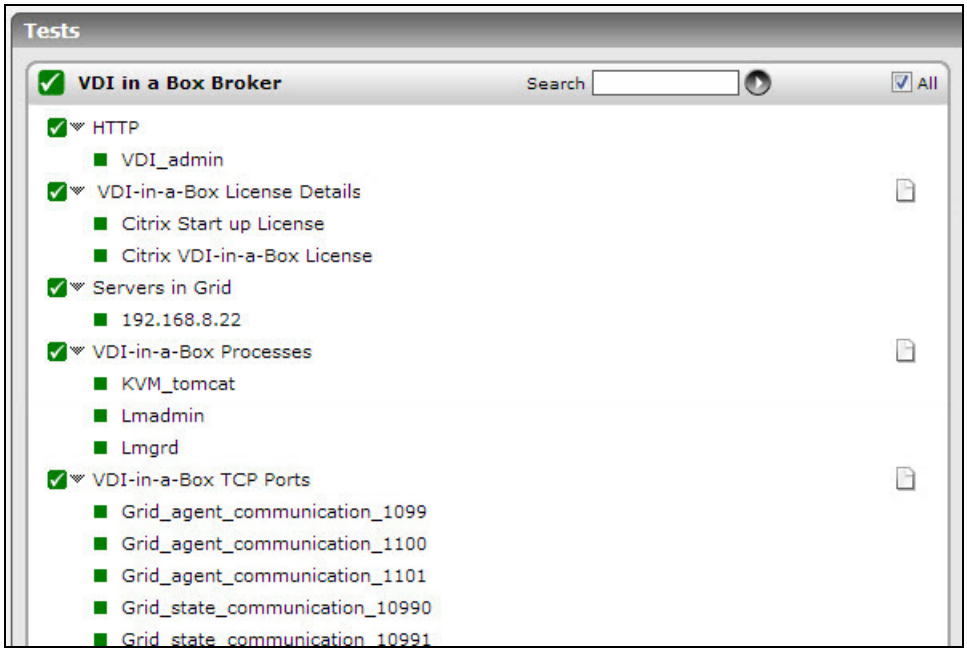

Figure 3.2: The tests mapped to the VDI In a Box Broker layer

3.1.1 HTTP Test

The vdimanager provides a web-based interface, the VDI-in-a-Box console, used to configure and manage servers running vdiManager, desktops, templates, images, users, and the grid, all at the grid level. In the VDI-in-a-Box console, the grid appears as one logical server running vdiManager. It is also possible to view the status and activity of each server individually when required, in the console. If the console is unreachable or takes too long to open, then the vdimanager cannot be configured to create desktops on-demand; user requests for desktops can hence not be serviced, resulting in a bad user experience with the entire virtual desktop service! The **HTTP** test seeks to avoid this unpleasant outcome by periodically checking the availability of the web console and its responsiveness to requests by emulating an HTTP request to the console from an external location.

Target of the test : A VDI-in-a-Box Manager

Agent deploying the test : An external agent; if you are running this test using the external agent on the eG manager box, then make sure that this external agent is able to communicate with the port

on which the target Webserver is listening. Alternatively, you can deploy the external agent that will be running this test on a host that can access the port on which the target console is listening.

Outputs of the test : One set of outputs for every URL being monitored.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which this test is to be configured.
Port	The port at which the specified host listens to.
URL	The web page being accessed. While multiple URLs (separated by commas) can be provided, each URL should be of the format *URL name:URL value*. *URL name* is a unique name assigned to the URL, and the *URL value* is the value of the URL. For example, a URL can be specified as *HomePage:http://192.168.10.12:7077/admin/*, where *HomePage* is the *URL name* and is the *URL value*. *http://192.168.10.12:7077/admin*
CookieFile	Whether any cookies being returned by the web server need to be saved locally and returned with subsequent requests.
ProxyHost	The host on which a web proxy server is running (in case a proxy server is to be used).
ProxyPort	The port number on which the web proxy server is listening.
ProxyUsername	The user name of the proxy server.
ProxyPassword	The password of the proxy server.
Confirm Password	Confirm the password by retyping it here.
Timeout	Here, specify the maximum duration (in seconds) for which the test will wait for a response from the server. The default Timeout period is 30 seconds.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Availability	This measurement indicates whether the server was able to respond successfully to the query made by the test.	Percent	Availability failures could be caused by several factors such as the web server process(es) being down, the web server being misconfigured, a network failure, etc. Temporary unavailability

Measurement	Description	Measurement Unit	Interpretation
			may also occur if the web server is overloaded. Availability is determined based on the response code returned by the server. A response code between 200 to 300 indicates that the server is available.
Total response time	This measurement indicates the time taken by the server to respond to the requests it receives.	Secs	Response time being high denotes a problem. Poor response times may be due to the server being overloaded or misconfigured. If the URL accessed involves the generation of dynamic content by the server, backend problems (e.g., an overload at the application server or a database failure) can also result in an increase in response time.
TCP connection availability	This measure indicates whether the test managed to establish a TCP connection to the server.	Percent	Failure to establish a TCP connection may imply that either the web server process is not up, or that the process is not operating correctly. In some cases of extreme overload, the failure to establish a TCP connection may be a transient condition. As the load subsides, the server may start functioning properly again.
TCP connect time	This measure quantifies the time for establishing a TCP connection to the web server host.	Secs	Typically, the TCP connection establishment must be very small (of the order of a few milliseconds). Since TCP connection establishment is handled at the OS-level, rather than by the application, an increase in this value signifies a system-level bottleneck on the host that supports the web server.
Server response time	This measure indicates the time period between when the connection was established and when the	Secs	While the total response time may depend on several factors, the server response time is typically, a very good indicator of a server bottleneck (e.g.,

Measurement	Description	Measurement Unit	Interpretation
	server sent back a HTTP response header to the client.		because all the available server threads or processes are in use).
Response code	The response code returned by the server for the simulated request	Number	A value between 200 and 300 indicates a good response. A 4xx value indicates a problem with the requested content (eg., page not found). A 5xx value indicates a server error.
Content length	The size of the content returned by the server	Kbytes	Typically the content length returned by the server for a specific URL should be the same across time. Any change in this metric may indicate the need for further investigation on the server side.
Content validity	This measure validates whether the server was successful in executing the request made to it.	Percent	A value of 100% indicates that the content returned by the test is valid. A value of 0% indicates that the content may not be valid. This capability for content validation is especially important for multi-tier web applications. For example, a user may not be able to login to the web site but the server may reply back with a valid HTML page where in the error message, say, "Invalid Login" is reported. In this case, the availability will be 100 % (since we got a valid HTML response). If the test is configured such that the content parameter should exclude the string "Invalid Login," in the above scenario content validity would have a value 0.

3.1.2 VDI-in-a-Box License Details Test

VDI-in-a-Box is licensed based on the count of concurrent users. By tracking the number of users who are simultaneously connected to virtual desktops via the vdimanager, you can easily figure out how the concurrent user licenses have been utilized.

The startup license on the other hand does not affect the VDI-in-a-Box license count. It is used to allow a Citrix product to communicate with the license server using a continuous open connection. Each time a Citrix product starts, it opens a connection to the license server by checking out the startup license. Every five minutes the license server and the products send a heartbeat message to each other to verify that they are mutually communicating.

This test tracks the usage of the Citrix Startup and VDI-in-a-Box licenses, and helps determine the following:

- Whether adequate Citrix startup licenses are available for supporting continuous connections to the Citrix License server;
- Whether sufficient concurrent user licenses are available or more licenses need to be purchased.

This way, you can ensure the uninterrupted use of the vdimanager and the virtual desktops it manages.

Target of the test : A VDI-in-a-Box Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for Citrix Startup licenses and Citrix CDI in-a-box licenses.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which this test is to be configured.
Port	The port at which the specified host listens to.
VDI in a Box Host	The IP address of the vdimanager that is being monitored.
VDI in a Box Username, VDI in a Box Password, and Confirm Password	The test needs to log on to the VDI-in-a-Box appliance to pull out the metrics of interest. For this purpose, you need to configure the test with the credentials of a user with login rights to the appliance. By default, the appliance supports a **root user** named *root* (with the default password root) and a *Read-only* user named **kvm** (with the default password **kaviza123**). If you prefer not to expose the credentials of the root user, then you can configure the VDI in a Box Username and VDI in a Box Password parameters with the credentials of the user **kvm**. Confirm the password of this user by retyping it in the Confirm Password text box. **Note:** The default password of user **kvm** is **kaviza123**. If this is changed, then make sure you configure the VDI in a box password parameter with the new password.

Parameters	Description
SSH Port	Specify the SSH port number of the vdimanager. By default, this is 22.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Licenses installed	Indicates the total number of licenses of this type that are currently installed.	Number	
Licenses in use	Indicates the number of licenses of this type that are currently in use.	Number	If the value of this measure is equal to that of Licenses installed, it indicates that all the licenses are utilized. If you have exhausted the startup licenses, then VDI-in-a-Box may not be able to communicate continuously with the license server and check out licenses, until additional licenses of this type are purchased. If no more VDI-in-box licenses are free, it implies that subsequent desktop requests from users will be rejected, until additional licenses of this type are purchased.
Available licenses	Indicates the number of licenses of this type that are currently unused.	Number	A high value is desired for this measure.
License utilization	Indicates the percentage of installed licenses that are currently in use.	Percent	A value close to 100% is indicative of rapid depletion of licenses; you may have to purchase additional licenses to continue using the product. If enough startup licenses are not available, then VDI-in-a-Box may not be able to communicate continuously with the license server and check out licenses until additional licenses of this type are purchased. If sufficient VDI-in-box licenses are not available, it implies that subsequent requests for desktops will be rejected by the

Measurement	Description	Measurement Unit	Interpretation
			vdimanager, until additional licenses of this type are purchased.

3.1.3 VDI-in-a-Box Processes Test

The vdimanager will not be available for use, if any/all of the following processes stop executing:

- The **KVM_tomcat** process
- **Lmadmin**, which is the Citrix license server manager process
- **Lmgrd**, which is the FLEXlm license manager daemon that handles the initial point of contact with the VDI-in-a-Box appliance, and processes all licensing requests

The **VDI-in-a-Box Processes** test has been pre-configured to monitor and report on the status (whether running or not) of the above-mentioned processes and their resource usage. If users are not able to connect to the vdimanager or complain of a slowdown when working with the vdimanager, you can use this test to figure out if one/all of these critical processes are at fault.

Target of the test : A VDI-in-a-Box Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each Process that is configured.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which this test is to be configured.
Port	The port at which the specified host listens to.
VDI in a Box Host	The IP address of the vdimanager that is being monitored.
VDI in a Box Username, VDI in a Box Password, and Confirm Password	The test needs to log on to the VDI-in-a-Box appliance to pull out the metrics of interest. For this purpose, you need to configure the test with the credentials of a user with login rights to the appliance. By default, the appliance supports a **root user** named *root* (with the default password root) and a *Read-only* user named **kvm** (with the default password **kaviza123**). If you prefer not to expose the credentials of the root user, then you can configure the VDI in a Box Username and VDI in a Box Password

Parameters	Description
	parameters with the credentials of the user **kvm**. Confirm the password of this user by retyping it in the Confirm Password text box. **Note:** The default password of user **kvm** is **kaviza123**. If this is changed, then make sure you configure the VDI in a box password parameter with the new password.
SSH Port	Specify the SSH port number of the vdimanager. By default, this is *22*.
Process	By default, this parameter is configured to monitor the *KVM_tomcat, lmadmin,* and *lmgrd* processes. If required, you can add more processes for monitoring by specifying these processes as a comma separated list of *processNames:processPattern* pairs. *processName* is a string that will be used for display purposes only. *processPattern* is an expression of the form - **expr** or *expr* or **expr* or *expr** or **expr1*expr2**... or *expr1*expr2*, etc. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. The pattern(s) used vary from one application to another and must be configured per application. For example, to monitor the **lmgrd** process, your specification can be, **Lmgrd:*CITRIX*vdimgr*lmgrd***, where * denotes zero or more characters. Other special characters such as slashes (\) can also be used while defining the process pattern. **Note:** The Process parameter supports process patterns containing the ~ character.
User	By default, this parameter has a value "*none*"; this means that the test monitors all processes that match the configured patterns, regardless of the user executing them. If you want the test to monitor the processes for specific users alone, then, specify a comma-separated list of users to be monitored in the user text box. For instance: *john,elvis,sydney* If multiple processes are configured for monitoring and multiple users are also configured, then the test will check whether the first process is run by the first user, the second process by the second user, and so on. For instance, if the processes configured are *java:java.exe,apache:*httpd** and the users configured are john,elvis, then the test will check whether user john is running the process java, and user elvis is running the process apache. Similarly, if multiple processes are configured, but a single user alone is configured, then the test will check whether the specified user runs each of the configured processes. However, if you want to check whether a single process, say java.exe, is run by multiple users - say, james and jane - then, you have to do the following: • Your User specification should be: *james,jane* (if the target host is a Unix host

Parameters	Description
	• Your Process configuration should be: *Process1:java.exe*,*Process2:java.exe*. The number of processes in this case should match the number of users. Such a configuration will ensure that the test checks for the *java.exe* process for both the users, *james* and *jane*.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Processes running	Indicates the number of instances of a process(es) currently executing on a host.	Number	This value indicates if too many or too few processes corresponding to an application are executing on the host.
CPU utilization	Indicates the percentage of CPU used by executing process(es) corresponding to the pattern specified.	Percent	A very high value could indicate that processes corresponding to the specified pattern are consuming excessive CPU resources.
Memory utilization	For one or more processes corresponding to a specified set of patterns, this value represents the ratio of the resident set size of the processes to the physical memory of the host system, expressed as a percentage.	Percent	A sudden increase in memory utilization for a process(es) may be indicative of memory leaks in the application.

3.1.4 VDI-in-a-Box TCP Ports Test

This test periodically checks and reports the availability and responsiveness of configured TCP ports on the vdimanager appliance. If users complain of delays when accessing the vdimanager or denial of access to the vdimanager, you can use this test to determine the reason for the anomaly – is it because, critical TCP ports on the vdimanager are currently unavailable? or is it due to the poor responsiveness of these ports to connection requests?

Target of the test : A VDI-in-a-Box Manager

Agent deploying the test : An external agent

Outputs of the test : One set of results for every configured port name.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which this test is to be configured.
Port	The port at which the specified host listens to.
VDI in a Box Host	The IP address of the vdimanager that is being monitored.
TargetPorts	Specify either a comma-separated list of port numbers that are to be tested (eg., 80,7077,1521), or a comma-separated list of *port name:port number* pairs that are to be tested (eg., VDI_manager:443,SSH:22). In the latter case, the port name will be displayed in the monitor interface. Alternatively, this parameter can take a comma-separated list of *port name:IP address:port number* pairs that are to be tested, so as to enable the test to try and connect to Tcp ports on multiple IP addresses. For example, *mysql:192.168.0.102:1433,egwebsite:209.15.165.127:80*.
Timeout	Here, specify the maximum duration (in seconds) for which the test will wait for a response from the server. The default Timeout period is 60 seconds.
IsPassive	If the value chosen is **Yes**, then the server under consideration is a passive server in a cluster. No alerts will be generated if the server is not running. Measures will be reported as “Not applicable’ by the agent if the server is not up.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Availability	Whether the TCP connection is available	Percent	An availability problem can be caused by different factors – e.g., the server process may not be up, a network problem may exist, or there could be a configuration problem with the DNS server.
Response time	Time taken (in seconds) by the server to respond to a request.	Secs	An increase in response time can be caused by several factors such as a server bottleneck, a configuration problem with the DNS server, a network problem, etc.

3.1.5 Servers in Grid Test

The grid unites servers running vdiManager, allowing load balancing and ensuring high availability of virtual machines on the servers. Servers are organized into a grid, and all the servers in a grid are functionally identical, so if one dies any of the others act in its place. For each server in the monitored grid, this test reports the current status of that server, the server configuration, the status of the virtual desktops on the server, and the disk space utilization on the server, so that:

- Inactive servers in the grid can be quickly identified;
- Servers experiencing a space crunch can be isolated;
- The number of desktops in various states of usage/activity can be ascertained

Target of the test : A VDI-in-a-Box Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each server in the target grid.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which this test is to be configured.
Port	The port at which the specified host listens to.
VDI in a Box Host	The IP address of the vdimanager that is being monitored.
VDI in a Box Username, VDI in a Box Password, and Confirm Password	The test needs to log on to the VDI-in-a-Box appliance to pull out the metrics of interest. For this purpose, you need to configure the test with the credentials of a user with login rights to the appliance. By default, the appliance supports a **root user** named *root* (with the default password root) and a *Read-only* user named **kvm** (with the default password **kaviza123**). If you prefer not to expose the credentials of the root user, then you can configure the VDI in a Box Username and VDI in a Box Password parameters with the credentials of the user **kvm**. Confirm the password of this user by retyping it in the Confirm Password text box. **Note:** The default password of user **kvm** is **kaviza123**. If this is changed, then make sure you configure the VDI in a box password parameter with the new password.
SSH Port	Specify the SSH port number of the vdimanager. By default, this is *22*.

Measurements made by the tesr

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Server status</td><td>Indicates the current status of this server.</td><td>Number</td><td>If the server is an active server in the grid, this measure will report the value Activated. If the server is an inactive server in the grid, this measure will report the value Deactivated.

The values that this measure can take and their corresponding numeric values are discussed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Activated</td><td>1</td></tr><tr><td>Deactivated</td><td>0</td></tr><tr><td>Quiescing</td><td>2</td></tr><tr><td>Deactivating</td><td>3</td></tr><tr><td>Shutting Down</td><td>4</td></tr><tr><td>Shutdown</td><td>5</td></tr><tr><td>Missing</td><td>6</td></tr><tr><td>Broken</td><td>7</td></tr></table>
Note:

By default, this measure reports the above-mentioned Measure Values while indicating the current status of a server. However, in the graph of this measure, the status will be represented using the corresponding numeric equivalents only.</td></tr>
<tr><td>Desktops in use</td><td>Indicates the number of desktops on this server that are currently in use.</td><td>Number</td><td></td></tr>
<tr><td>Desktops preserved for subsequent use</td><td>Indicates the number of desktops on this server that have been preserved for subsequent use.</td><td>Number</td><td>This refers to the number of desktops that are ‘On Hold’.</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Desktops pre-started for instant login</td><td>Indicates the number of desktops on this server that are already prestarted for instant login.</td><td>Number</td><td rowspan="2">Pre-started desktops are in a powered-on state and at the logon prompt, ready for use. Pre-starting desktops eliminates the need for users to wait for virtual desktops to start.</td></tr>
<tr><td>Desktops being pre-started for instant login</td><td>Indicates the number of desktops on this server that are being pre-started for instant login.</td><td>Number</td></tr>
<tr><td>Server type</td><td>Indicates the type of this server.</td><td></td><td>The values that this measure can take and their corresponding numeric values are discussed in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>VMware</td><td>1</td></tr><tr><td>XenServer</td><td>2</td></tr><tr><td>Hyper-V</td><td>3</td></tr></table>
Note:
By default, this measure reports the above-mentioned Measure Values while indicating the server type. However, in the graph of this measure, the type will be represented using the corresponding numeric equivalents only.
If you have added a server of a different type, say XenServer, to a grid of servers of type VMware, then this measure will reveal this inconsistency.</td></tr>
<tr><td>Total RAM size</td><td>Indicates the total RAM size of this server.</td><td>MB</td><td></td></tr>
<tr><td>CPU cores</td><td>Indicates the number of CPU cores of this server.</td><td>Number</td><td></td></tr>
<tr><td>Total CPU</td><td>Indicates the total CPU cycles available across all CPU cores in the server.</td><td>GHz</td><td></td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
Total capacity	Indicates the total capacity of the logical storage device attached to this server.	GB	
Free space	Indicates the amount of unused space in the logical storage device of this server.	GB	A high value is desired for this measure.
Used space	Indicates the amount of space that is currently in use on the logical storage device of this server.	GB	A very high value is a cause for concern, as it indicates that the server is running short of storage space. Without sufficient space, critical read/write operations cannot be performed by the server. You may then have to clear some space in the logical storage devices or allocate more space to these devices to avert the problem.
Space utilization	Indicates the percentage of space in the logical storage device of this server that is currently in use.		

3.2 The Virtual Desktops Layer

To know the current status and usage of templates, images, and desktops configured on the vdimanager, use the tests mapped to this layer.

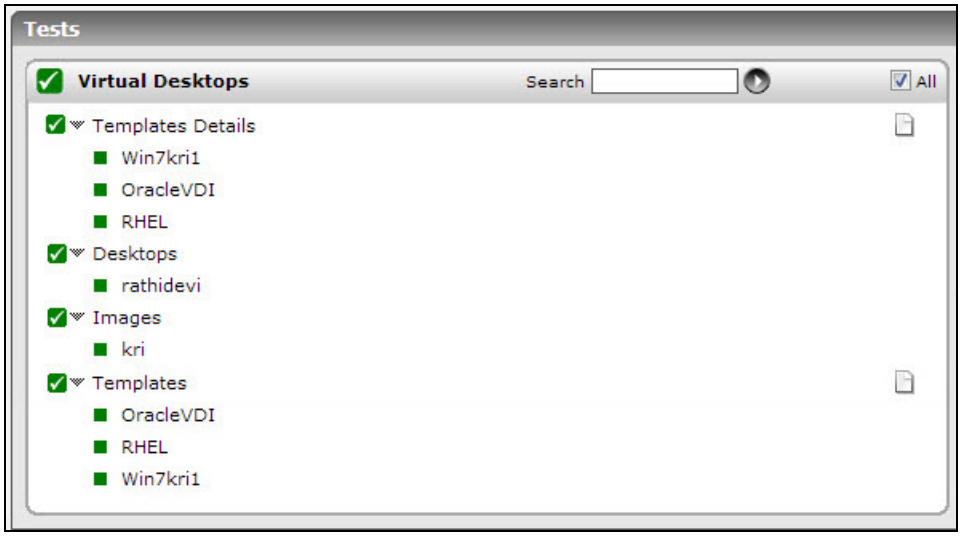

Figure 3.3: The tests mapped to the Virtual Desktops layer

3.2.1 Templates Test

Templates are the molds from which desktops are created. You can use templates to create uniform virtual desktops that meet your specifications. Templates consist of an image and policies. The image contains the operating systems and applications that run on the desktop. Policies, which you set while creating the template, are characteristics such as how many desktops to create and how much RAM to allocate to the desktops. The Templates test allows you a sneak peek at the configuration of each template created using the vdimanager and the status of every template so created. This way, you can identify templates that are currently enabled and those that are not, and those templates from which desktops with the configuration you desire can be created.

Target of the test : A VDI-in-a-Box Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each server in the target grid.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which this test is to be configured.
Port	The port at which the specified host listens to.
VDI in a Box Host	The IP address of the vdimanager that is being monitored.
VDI in a Box Username, VDI in a Box Password, and Confirm Password	The test needs to log on to the VDI-in-a-Box appliance to pull out the metrics of interest. For this purpose, you need to configure the test with the credentials of a user with login rights to the appliance. By default, the appliance supports a **root user** named *root* (with the default password root) and a *Read-only* user named **kvm** (with the default password **kaviza123**). If you prefer not to expose the credentials of the root user, then you can configure the VDI in a Box Username and VDI in a Box Password parameters with the credentials of the user **kvm**. Confirm the password of this user by retyping it in the Confirm Password text box. **Note:** The default password of user **kvm** is **kaviza123**. If this is changed, then make sure you configure the VDI in a box password parameter with the new password.
SSH Port	Specify the SSH port number of the vdimanager. By default, this is *22*.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an

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

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Status	Indicates the current status of this template.		The values that this measure can take and their corresponding numeric values are discussed in the table below: Measure Value: Enabled — Numeric Value: 1 Measure Value: Disabled — Numeric Value: 0 **Note:** By default, this measure reports the above-mentioned **Measure Value**s while indicating the current status of a template. However, in the graph of this measure, the status will be represented using the corresponding numeric equivalents only.
Is default template?	Indicates whether this template is the default template or not.		A desktop from the Default template is provisioned to a user if the user is not assigned to any template. The values that this measure can take and their corresponding numeric values are discussed in the table below:

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Yes</td><td>1</td></tr>
<tr><td>No</td><td>0</td></tr>
</table>
Note:
By default, this measure reports the above-mentioned Measure Values while indicating whether a template is the default template or not. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.
</td></tr>
<tr><td></td><td></td><td></td><td>To view the complete details of the monitored template, use the detailed diagnosis of this measure. This will reveal the policies and properties set for the desktops that will be created from the target template.</td></tr>
<tr><td>Refresh policy</td><td>Indicates the refresh policy configured for the desktops that will be created from this template.</td><td></td><td>The values that this measure can take and their corresponding numeric values are discussed in the table below:
<table>
<tr><th>Measure Value</th><th>Description</th><th>Numeric Value</th></tr>
<tr><td>On logout</td><td>Ensures that users get a pristine desktop each time they login by deleting the desktop as soon as the user logs out</td><td>1</td></tr>
<tr><td>Scheduled</td><td>Refreshes all "On Hold" or "On Hold" and "In Use" desktops at the scheduled refresh time on a daily, weekly, or</td><td>2</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Description</th><th>Numeric Value</th></tr>
<tr><td></td><td>monthly basis.</td><td></td></tr>
<tr><td>Scheduled or on logout</td><td>Refreshes desktops at the scheduled refresh time as well as when the user logs out</td><td>3</td></tr>
<tr><td>Manual</td><td>Does not refresh any desktop unless the administrator does so manually</td><td>4</td></tr>
<tr><td>Personal</td><td>Desktops are refreshed only when the base image is updated, and the refresh reintegrates the existing personal disk with the new base image</td><td>5</td></tr>
</table>
Note:
By default, this measure reports the above-mentioned Measure Values while indicating the refresh policy of a template. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.
</td></tr>
<tr><td>Memory</td><td>Indicates the amount of RAM allocated to the desktops created from this template.</td><td>MB</td><td></td></tr>
<tr><td>Virtual CPU</td><td>Indicates the number of virtual cores allocated to the desktops created from this template.</td><td>Number</td><td></td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
Maximum desktops	Indicates the maximum number of desktops that can be generated from this template.	Number	
Pre-started desktops	Indicates the number of desktops created from this template that should be pre-started and ready for login.	Number	Pre-started desktops are in a powered-on state and at the logon prompt, ready for use. Pre-starting desktops eliminates the need for users to wait for virtual desktops to start.

The detailed diagnosis of the *Is default template?* measure reveals the policies and properties set for the desktops that will be created from a target template. This includes information such as the policy type, the desktop type, and whether serial ports, disk drives, printers, and smart cards are to be enabled for the desktops or not.

Shows the details templates

TIME	DESCRIPTION	FAST REFRESH	DESKTOP TYPE	POLICY TYPE	SERIAL PORTS	DISK DRIVES	PRINTERS	SMART CARDS
Jan 04, 2013 18:50:34								
	VdiTesting	Enabled	Pooled	Do not reassign desktops "On Hold" to new users	Not Enabled	Enabled	Enabled	Not Enabled

Figure 3.4: The detailed diagnosis of the Is default template? Measure

3.2.2 Templates Details Test

Templates are the molds from which desktops are created. You can use templates to create uniform virtual desktops that meet your specifications. Templates consist of an image and policies. The image contains the operating systems and applications that run on the desktop. Policies, which you set while creating the template, are characteristics such as how many desktops to create and how much RAM to allocate to the desktops. Using the **Templates Details** test, you can determine the current status of the desktops created using each template configured using the vdimanager. This way, you can figure out how many desktops are currently in use and how many are broken.

Target of the test : A VDI-in-a-Box Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each template.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which this test is to be configured.
Port	The port at which the specified host listens to.
VDI in a Box Host	The IP address of the vdimanager that is being monitored.
VDI in a Box Username, VDI in a Box Password, and Confirm Password	The test needs to log on to the VDI-in-a-Box appliance to pull out the metrics of interest. For this purpose, you need to configure the test with the credentials of a user with login rights to the appliance. By default, the appliance supports a **root user** named *root* (with the default password root) and a *Read-only* user named **kvm** (with the default password **kaviza123**). If you prefer not to expose the credentials of the root user, then you can configure the VDI in a Box Username and VDI in a Box Password parameters with the credentials of the user **kvm**. Confirm the password of this user by retyping it in the Confirm Password text box. **Note:** The default password of user **kvm** is **kaviza123**. If this is changed, then make sure you configure the VDI in a box password parameter with the new password.
SSH Port	Specify the SSH port number of the vdimanager. By default, this is *22*.
Detailed Diagnosis	To make diagnosis more efficient and accurate, eG Enterprise embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormalfrequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Desktops in use	Indicates the number of desktops created from this template that are currently in use.	Number	This refers to the number of desktops to which a user is currently logged into. The detailed diagnosis of this measure will list the names of desktops that are currently in use.
Desktops not in use	Indicates the number of desktops created from this template that are currently not in use.	Number	
Desktops utilization	Indicates the percentage of desktops created from this template that are currently in use.	Percent	A high value for this measure indicates that too many desktops are in use. This in turn implies that very few desktops are currently available for usage.
Desktops pre-started for instant login	Indicates the number of desktops created from this template that are ready for login – i.e., the number of new desktops.	Number	
Desktops preserved for subsequent use	Indicates the number of desktops created from this template that were in use in the past, but the user has currently logged out.	Number	This refers to the number of desktops that are ‘on hold’. You can use the detailed diagnosis of this measure to know which desktops are on hold currently.
Broken desktops	Indicates the number of desktops created from this template that are broken.	Number	A desktop in an unknown state has the status "Broken" and is not available to users. Causes for broken desktops are as follows: • Desktops are getting stuck in startup within Windows due to missing or incorrect startup configuration, e.g., a bad product key or incorrectly specified domain

Measurement	Description	Measurement Unit	Interpretation
			credentials. This can be diagnosed by using a console client to directly look at the console of the broken desktop to see if it is waiting on user input. This is the most common reason desktops are broken, and this cause can be mitigated by thoroughly testing working desktops before saving them as templates. • The server is too heavily loaded to start new desktops. • The disk space on the datastore is exhausted. • The computer name allocation system has been exhausted. • The MAC address allocation system has been exhausted. • The template image has been lost. To know which desktops are broken, use the detailed diagnosis of this measure.
Desktops being pre-started for instant login	Indicates the number of desktops created from this template that are being pre-started for instant login.	Number	

To know which desktops are currently in use for a template, use the detailed diagnosis of the **Desktops in use** measure.

Shows the details of used desktops

TIME	USER NAME	TEMPLATE NAME	SERVER IP	VM NAME	DESKTOP ADDRESS	CLIENT ADDRESS	LOGIN TIME	STATUS
Jan 04, 2013 18:54:24								
	rathidevi	Win7kri1	192.168.8.22	Win7kri15	192.168.9.158	192.168.8.192	Jan 2, 2013 11:19 PM	Active

Figure 3.5: The detailed diagnosis of the Desktops in use measure

To know which desktops for a template are currently on hold, use the detailed diagnosis of the `Desktops preserved for subsequent use` measure.

Shows the details of holding desktops

TIME	USER NAME	TEMPLATE NAME	SERVER IP	VM NAME	DESKTOP ADDRESS	CLIENT ADDRESS	LOGIN TIME	STATUS
Jan 04, 2013 18:54:24								
	eguser	Win7kri1	192.168.8.22	Win7kri16	192.168.9.179	192.168.8.192	Jan 3, 2013 1:34 AM	Logging in...

Figure 3.6: The detailed diagnosis of the Desktops preserved for subsequent use measure

To know the broken desktops, use the detailed diagnosis of the *Broken desktops* measure.

Shows the details of broken desktops

TIME	USER NAME	TEMPLATE NAME	SERVER IP	VM NAME	DESKTOP ADDRESS	CLIENT ADDRESS	LOGIN TIME	STATUS
Jan 07, 2013 12:50:09								
	N/A	Win7kri1	192.168.8.22	Win7kri19	192.168.9.187	N/A	N/A	Join domain: Domain join failed with error code 1355 reason(if available) The specified domain either does not exist or could not be contacted

Figure 3.7: The detailed diagnosis of the Broken desktops measure

3.2.3 Desktops Test

If a user complains of issues when accessing or working with a desktop, administrators must be able to quickly zoom into that user's session to identify the desktop that the user is currently logged into and determine what is wrong with the user's session – is it inactive? Is it awaiting a refresh? Has the session logged out? The **Desktops** test provides administrators with this insight. This test auto-discovers all users who are currently logged into desktops and reports on the state of each user's session, so that problems (if any) with desktop sessions can be isolated and fixed.

Target of the test : A VDI-in-a-Box Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each user session.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which this test is to be configured.
Port	The port at which the specified host listens to.
VDI in a Box Host	The IP address of the vdimanager that is being monitored.
VDI in a Box Username, VDI in a Box Password, and Confirm Password	The test needs to log on to the VDI-in-a-Box appliance to pull out the metrics of interest. For this purpose, you need to configure the test with the credentials of a user with login rights to the appliance. By default, the appliance supports a **root user** named *root* (with the default password root) and a *Read-only* user named **kvm** (with the default password **kaviza123**). If you prefer not to expose the credentials of the root user, then you can configure the VDI in a Box Username and VDI in a Box Password parameters with the credentials of the user **kvm**. Confirm the password of this user by retyping it in the Confirm Password text box. **Note:** The default password of user **kvm** is **kaviza123**. If this is changed, then make sure you configure the VDI in a box password parameter with the new password.
SSH Port	Specify the SSH port number of the vdimanager. By default, this is *22*.
Detailed Diagnosis	To make diagnosis more efficient and accurate, eG Enterprise embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Status</td><td>Indicates the current status of this user's desktop session.</td><td></td><td>The values that this measure can take and their corresponding numeric values are discussed in the table below:
<table>
<tr><th>Measure Value</th><th>Description</th><th>Numeric Value</th></tr>
<tr><td>Active</td><td>User session is active. It is also current, i.e., the user session has started after the last desktop refresh cycle.</td><td>1</td></tr>
<tr><td>Pending refresh</td><td>User session is active. However, it is active beyond the last desktop refresh cycle. The user had logged in before the last desktop refresh cycle and has not logged out yet</td><td>2</td></tr>
<tr><td>On Hold</td><td>A user had logged in and logged out after the last desktop refresh cycle</td><td>3</td></tr>
<tr><td>In Use</td><td>The desktop session is displaying a login screen on a kiosk-style client</td><td>4</td></tr>
<tr><td>Logging in</td><td>The desktop session is logged out and ready to accept a new login on a kiosk- style</td><td>5</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Description</th><th>Numeric Value</th></tr>
<tr><td></td><td>client
The user is in the process of logging in.</td><td></td></tr>
<tr><td>Kiosk ready</td><td>The desktop session is displaying a login screen on a kiosk-style client</td><td>6</td></tr>
<tr><td>Kiosk</td><td>The desktop session is logged out and ready to accept a new login on a kiosk-style client</td><td>7</td></tr>
</table>
Note:
By default, this measure reports the above-mentioned Measure Values while indicating the session state. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.
To know which desktop the user is currently logged into, use the detailed diagnosis of this measure.
</td></tr>
</table>

Use the detailed diagnosis of the *Status* measure to know which desktop the user is logged into, from which client the user logged in, when he/she logged in, and the duration of access.

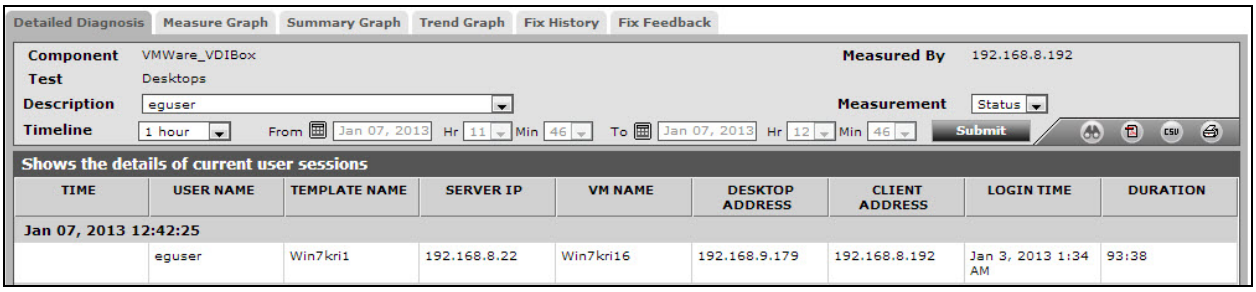

Figure 3.8: The detailed diagnosis of the Status measure of the Desktops test

3.2.4 Images Test

An image includes a desktop operating system (such as Windows 7 or Windows XP), a set of applications, and the VDI-in-a-Box Desktop Agent; without this agent, the vdimanager will not be able to receive updates on user sessions and desktop health. The vdimanager creates desktops from templates that consist of an image. A single image can be used by more than one template. Once an image is created, you can distribute that image to all the other servers in a grid. The **Images** test not only reports the distribution status (whether the image is successfully distributed or not) of an image, but also periodically checks the configuration of an image and reports whether the Desktop Agent is installed on the image or not. This way, the test promptly alerts you to distribution failures, points you to the reason why the vdimanager has not been receiving updates on user sessions and desktop health, and also leads you to the specific templates that have been affected by such incomplete (i.e., images without desktop agent) images.

Target of the test : A VDI-in-a-Box Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each user session

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The host for which the test is to be configured.
Port	The port number at which the specified Host listens to.
Vdi in a box host	The IP address of the vdimanager being monitored
Vdi in a box Username, Vdi in a	The test needs to log on to the VDI-in-a-Box appliance to pull out the metrics of interest. For this purpose, you need to configure the test with the credentials of a user

Parameter	Description
box Password, and Confirm Password	with login rights to the appliance. By default, the appliance supports a root user named root (with the default password root) and a Read-only user named *kvm* (with the default password *kaviza123*). If you prefer not to expose the credentials of the root user, then you can configure the Vdi in a box Username and Vdi in a box Password parameters with the credentials of the user *kvm*. Confirm the password of this user by retyping it in the Confirm Password text box. **Note:** The default password of user kvm is *kaviza123*. If this is changed, then make sure you configure the Vdi in a box Password parameter with the new password.
SSH Port	Specify the SSH port number of the vdimanager. By default, this is 22.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Status	Indicates the current status of this image.		The values that this measure can take and their corresponding numeric values are discussed in the table below: **Measure Value** / **Description** / **Numeric Value** Green / Image is distributed throughout / 1

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Description</th><th>Numeric Value</th></tr>
<tr><td></td><td>the grid or fully distributed.</td><td></td></tr>
<tr><td>Yellow</td><td>Image distribution is ongoing</td><td>2</td></tr>
<tr><td>Red</td><td>Image distribution is complete.</td><td>3</td></tr>
</table>
Note:
By default, this measure reports the above-mentioned Measure Values while indicating the distribution state of an image. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.
</td></tr>
<tr><td>Image distribution</td><td>Indicates whether this image has been distributed or is being distributed to the servers in a grid.</td><td></td><td>
The values that this measure can take and their corresponding numeric values are discussed in the table below:
<table>
<tr><th>Measure Value</th><th>Description</th><th>Numeric Value</th></tr>
<tr><td>Published</td><td>The image is fully distributed throughout the grid</td><td>1</td></tr>
<tr><td>Distributing</td><td>The image is being distributed throughout the grid</td><td>2</td></tr>
</table>
Note:
By default, this measure reports the above-mentioned Measure Values while
</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>indicating the distribution state of an image. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.</td></tr>
<tr><td>Total templates using this image</td><td>Indicates the number of templates using this image.</td><td>Number</td><td>If the desktop agent is not installed on an image, then using this measure, you can easily assess the extent of the damage – this is because, this measure reports the total number of templates that will be affected by the non-availability of the desktop agent in the underlying image. To know which templates will be affected, use the detailed diagnosis of this measure.</td></tr>
<tr><td>Is HDX remote protocol enabled?</td><td>Indicates whether the HDX protocol is enabled for this image or not.</td><td></td><td>HDX is the default protocol that user devices use to communicate with virtual desktops provisioned by the VDI-in-a-Box appliance. HDX technology is a set of capabilities that deliver a “high definition” experience to end users of any application, on any device and over any network. This protocol provides multi-media support while using less bandwidth. It is suitable for remote access over a WAN.

The values that this measure can take and their corresponding numeric values are discussed in the table below:

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

By default, this measure reports the above-mentioned Measure Values while</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>indicating the state of the HDX protocol. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.</td></tr>
<tr><td>Is desktop agent installed?</td><td>Indicates whether the Desktop Agent has been installed on this image or not.</td><td></td><td>If the value of this measure is <code>No</code>, it indicates that the desktop agent has not been installed on the image. This implies that the vdimanager will not be able to receive any communication related to user connections or desktop health from the desktops created using this image. If the value of this measure is <code>Yes</code>, it indicates that the desktop agent has been installed.

The numeric values that correspond to the measure values discussed above are as follows:

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

By default, this measure reports the above-mentioned Measure Values while indicating the state of the HDX protocol. However, in the graph of this measure, the same will be represented using the corresponding numeric equivalents only.</td></tr>
<tr><td>Image size</td><td>Indicates the current size of this image.</td><td>MB</td><td></td></tr>
</table>

Use the detailed diagnosis of the *Total templates using the image* measure to know which templates are using a particular image.

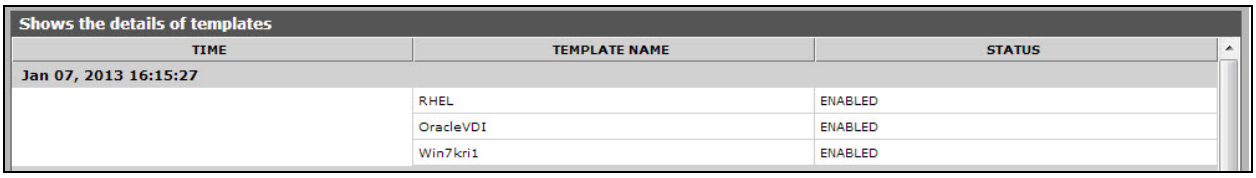

Figure 3.9: The detailed diagnosis of the Total templates using the image measure

To know the current configuration of an image, use the detailed diagnosis of the *Is desktop agent installed?* measure.

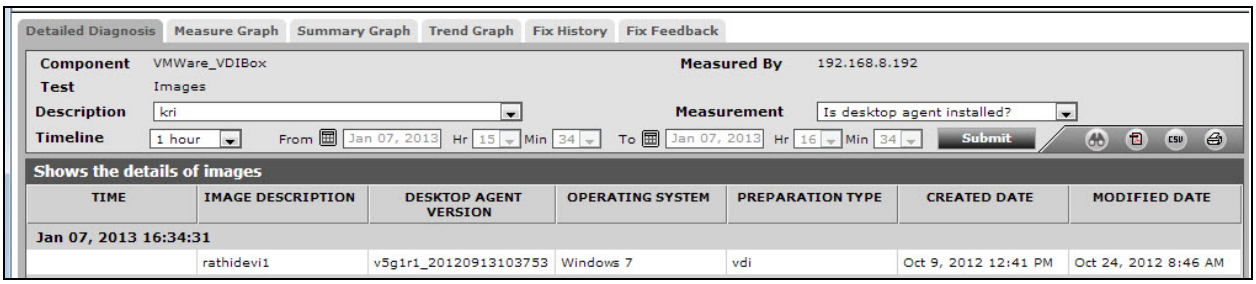

Figure 3.10: The detailed diagnosis of the Is desktop agent installed? measure

3.2.5 Virtual Desktop User Logon Experience Test

When a user clicks on an icon to launch an application delivered by VDI server, you will see a description of the processes that are taking place to launch the requested application. Any problem experienced in the launch process - within the VDI infrastructure or the supporting infrastructure components (e.g., the Active Directory server, the License server, the Web Interface server, the Profile server, etc.) - can result in a poor user experience. Inevitably, these issues result in service desk calls and complaints that “VDI is slow.” Diagnosing login problems has traditionally been a difficult, time-consuming, manual process due to the large number of steps involved, many of which are external to the VDI server. Critical operations such as user login script execution and profile loading can significantly elongate the overall login experience for a VDI user. The key to resolving user experience issues therefore, lies in tracking each user session with the VDI server, ascertaining the time spent by the session at each step of the logon process - be it within or external to the VDI server - and accurately identifying at what step of the logon process, the slowdown occurred - when user credentials are obtained? When credentials are validated? During profile loading? During login script execution? When mapping drives or creating printers? This is exactly what the **Virtual Desktop User Logon Experience** does.

This test captures every user login to the VDI server, provides visibility into each individual step involved in the login process of that user, and the time required for each step to finish. For this purpose, the test categorizes its metrics into client start-up metrics and server start-up metrics.

The *client start-up* metrics are concerned with timing the operations that occur from the point when the user requests an application, e.g., by clicking an icon, to the point at which an instance of the ICA client has finished opening a connection to Presentation Server. While connection- brokering mechanisms, such as Web Interface for VDI server or Program Neighborhood Agent, involve components that are not on the physical client device, the tasks these systems perform have a direct impact on the performance of the connection start-up and are recorded as part of the client-side process.

The *server start-up* metrics are concerned with timing the operations that occur when creating a new session on the VDI server. This includes user authentication, client device mapping, profile loading, login scripts execution, and finally, starting the user's application (in the case of a desktop this will be explorer.exe). If a session already exists and a new application is being started through session sharing, only the application start stage will be considered as part of server start-up.

The total time required for these two areas to complete will directly affect the experience of the user while waiting for their application to become ready. This way, the test pinpoints those users whose logins were slow, and also reveals the specific causes of slow logins. This visibility enhances the performance of applications delivered by VDI server, resulting in a better user experience.

Target of the test : A VDI-in-a-Box Manager

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for every *user_on_VM*.

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which this test is to be configured.
Port	The port at which the specified host listens to. By default, this is *NULL*.
Virtual Center	If the eG manager had discovered the target ESX server by connecting to vCenter, then the IP address of the vCenter server used for discovering this ESX server would be automatically displayed against the Virtual Center parameter; similarly, the ESX User and ESX Password text boxes will be automatically populated with the vCenter user credentials, using which ESX discovery was performed. If this ESX server has not been discovered using vCenter, but you still want to monitor the ESX server via vCenter, then select the IP address of the vCenter host that you wish to use for monitoring the ESX server from the Virtual Center list. By default, this

Parameters	Description
	list is populated with the IP address of all vCenter hosts that were added to the eG Enterprise system at the time of discovery. Upon selection, the ESX User and ESX Password that were pre-configured for that vCenter server will be automatically displayed against the respective text boxes. On the other hand, if the IP address of the vCenter server of interest to you is not available in the list, then, you can add the details of the vCenter server on-the-fly, by selecting the Other option from the VIRTUAL CENTER list. This will invoke the ADD VCENTER SERVER DETAILS page. Refer to Section 3.1.1.1.1 to know how to add a vCenter server using this page. Once the vCenter server is added, its IP address, ESX User, and ESX Password will be displayed against the corresponding text boxes. On the other hand, if you want the eG agent to behave in the default manner -i.e., communicate with each ESX server for monitoring it - then set the Virtual Center parameter to '*none*'. In this case, the ESX User and ESX Password parameters can be configured with the credentials of a user who has at least 'Read-only' privileges to the target ESX server.
Inside View Using	By default, this test communicates with every VM remotely and extracts "inside view" metrics. Therefore, by default, the inside view using flag is set to Remote connection to VM (Windows). Typically, to establish this remote connection with Windows VMs in particular, eG Enterprise requires that the eG agent be configured with domain administrator privileges. In high-security environments, where the IT staff might have reservations about exposing the credentials of their domain administrators, this approach to extracting "inside view" metrics might not be preferred. In such environments therefore, eG Enterprise provides administrators the option to deploy a piece of software called the eG VM Agent on every Windows VM; this VM agent allows the eG agent to collect "inside view" metrics from the Windows VMs without domain administrator rights. Refer to 1.3.2 for more details on the eG VM Agent. To ensure that the "inside view" of Windows VMs is obtained using the eG VM Agent, set the inside view using flag to eG VM Agent (Windows). Once this is done, you can set the Domain, Admin User, and Admin Password parameters to *none*.
Domain, Admin User, Admin Password, and Confirm Password	By default, this test connects to each virtual guest remotely and attempts to collect "inside view" metrics. In order to obtain a remote connection, the test must be configured with user privileges that allow remote communication with the virtual guests. The first step towards this is to specify the Domain within which the virtual guests reside. The Admin User and Admin Password will change according to the domain specification. Discussed below are the different values that the domain parameter can take, and how they impact the Admin User and Admin Password specifications:

Parameters	Description
	• **If the VMs belong to a single domain**: If the guests belong to a specific domain, then specify the name of that domain against the domain parameter. In this case, any administrative user in that domain will have remote access to all the virtual guests. Therefore, an administrator account in the given domain can be provided in the Admin User field and the corresponding password in the Admin Password field. Confirm the password by retyping it in the Confirm Password text box. • **If the VMs belong to different domains**: In this case, you might want to provide multiple domain names. If this is done, then, to access the guests in every configured domain, the test should be configured with the required user privileges; this implies that along with multiple Domain names, multiple Admin User names and Admin Passwords would also have to be provided. To help administrators provide these user details quickly and easily, the eG administrative interface embeds a special configuration page. To access this page, simply click on the **Click here** hyperlink that appears just above the parameters of this test in the test configuration page. To know how to use the special page, refer to Section 2.6.1.1 of this document. • **If the guests do not belong to any domain** : In this case, specify "none" in the Domain field, and specify a local administrator account name in the Admin User below. Prior to this, you need to ensure that the same local administrator account is available or is explicitly created on each of the virtual machines to be monitored. Then, proceed to provide the password of the Admin User against Admin Password, and confirm the password by retyping it in the Confirm Password text box. If key-based authentication is implemented between the eG agent and the SSH daemon of a Linux guest, then, in the Admin User text box, enter the name of the user whose <USER_HOME_DIR> (on that Linux guest) contains a .ssh directory with the public key file named authorized_keys. The Admin Password in this case will be the passphrase of the public key; the default public key file that is bundled with the eG agent takes the password eginnovations. Specify this as the Admin Password if you are using the default private/public key pair that is bundled with the eG agent to implement key-based authentication. On the other hand, if you are generating a new public/private key pair for this purpose, then use the passphrase that you provide while generating the pair. For the detailed procedure on Implementing Key-based Authentication refer to the *Monitoring Xen Servers* document.

Parameters	Description
	• **If the inside view using flag is set to 'eG VM Agent (Windows)'**: On the other hand, if the Inside View Using flag is set to **eG VM Agent** (Windows) , then it implies that the inside view can be obtained without *Domain Administrator* privileges. Therefore, set the Domain, Admin User, and Admin Password parameters to *none*.
Ignore VMs Inside View	Administrators of some high security Hyper-V VDI environments might not have permissions to internally monitor one/more VMs. The eG agent can be configured to not obtain the 'inside view' of such 'inaccessible' VMs using the Ignore VMs Inside View parameter. Against this parameter, you can provide a comma-separated list of VM names, or VM name patterns, for which the inside view need not be obtained. For instance, you ignore vms inside view specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to none indicating that the eG agent obtains the inside view of all VMs on a virtual host by default. **Note:** While performing VM discovery, the eG agent will not discover the operating system of the VMs configured in the Ignore VMs Inside View text box.
Exclude VMs	Administrators of some virtualized environments may not want to monitor some of their less-critical VMs - for instance, VM templates - both from 'outside' and from 'inside'. The eG agent in this case can be configured to completely exclude such VMs from its monitoring purview. To achieve this, provide a comma-separated list of VMs to be excluded from monitoring in the Exclude VMs text box. Instead of VMs, VM name patterns can also be provided here in a comma-separated list. For example, your Exclude VMs specification can be: **xp,*lin*,win*,vista*. Here, the * (asterisk) is used to denote leading and trailing spaces (as the case may be). By default, this parameter is set to *none* indicating that the eG agent obtains the inside and outside views of all VMs on a virtual host by default. By providing a comma-separated list of VMs/VM name patterns in the Exclude VMs text box, you can make sure the eG agent stops collecting 'inside' and 'outside' view metrics for a configured set of VMs.
Ignore WINNT	By default, the eG agent does not support the inside view for VMs executing on Windows NT operating systems. Accordingly, the Ignore WINNT flag is set to **Yes**, by default.
Report By User	While monitoring a Hyper-V VDI, the Report By User flag is set to **No** by default, indicating that by default, the guest operating systems on the Hyper-V VDI are identified using the hostname specified in the operating system. On the other hand, while monitoring a Hyper-V VDI, this flag is set to **Yes** by default; this implies that in

Parameters	Description
	case of the Hyper-V VDI model, by default, the desktops will be identified using the login of the user who is accessing them. In other words, in VDI environments, this test will, by default, report measures for every u*sername_on_virtualmachinename*.
Report Powered OS	**This flag becomes relevant only if the Report By User flag is set to 'Yes'**. If the Report Powered OS Flag is set to **Yes** (which is the default setting), then this test will report measures for even those VMs that do not have any users logged in currently. Such guests will be identified by their virtualmachine name and not by the *username_on_virtualmachinename*. On the other hand, if the Report Powered OS flag is set to **No**, then this test will not report measures for those VMs to which no users are logged in currently.
Report By Managertime	By default, this flag is set to **Yes**. This indicates that the user login time displayed in the **DETAILED DIAGNOSIS** page for this test and in the Thin Client reports will be based on the eG manager's time zone by default. Set this flag to **No** if you want the login times displayed in the **DETAILED DIAGNOSIS** page for this test and in the Thin Client reports to be based on the virtual server's local time.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
User sessions	Indicates the number of sessions currently open for this user on the VDI server.	Number	Compare the value of this measure across users to know which user has the maximum number of open sessions. In the event of an overload, this will point you to that user who is contributing the most to the workload of the VDI server. Use the detailed diagnosis of this measure to view the complete details of each user session. Such details includes the name and IP address of the client from which every session was launched, when session creation started, and when it ended. With the help of this information, administrators can quickly pinpoint those sessions that may have taken too long to be created.
Logon duration	Indicates the time taken by this user to complete the logon process.	Secs	Compare the value of this measure across users to know which user took the longest to logon to the VDI server. To know what is causing this 'slow login', compare the values reported by all the other 'duration' measures of this test for that user. This will quickly lead you to where that user's logon process is spending the maximum time.
Profile load duration	Indicates the time taken to load this user's profile.	Secs	Compare the value of this measure across users to know which user's profile took the longest to load. If that user's Logon duration is high, you may want to compare the value of this measure with that of the other 'duration' measures reported for this user to figure out if a delay in profile loading is what is really ailing that user's logon experience.

Measurement	Description	Measurement Unit	Interpretation
			One of the common reasons for high profile load time is the large size of the user profile. Use the **User Profile** test to know if this user's profile is indeed exceeding the configured size limit. If so, use the **User Profile** test again to know which specific files in that user's profile are contributing to its large size, and check if such files can be deleted to reduce profile size, and consequently, the logon duration. Moreover, this measure reports the average time taken for loading a user's profile across all the sessions of that user. To know the profile load time per user session, use the detailed diagnosis of this measure. This will accurately pinpoint the session in which the profile took the longest to load.
Group policy processing duration	Indicates the time taken by this user's session to process group policies.	Secs	Compare the value of this measure across users to know which user's sessions took the longest time to process group policies. If that user's Logon duration is high, you may want to compare the value of this measure with that of the other 'duration' measures reported for this user to figure out if a delay in group policy processing is what is really ailing that user's logon experience. Logon performance improves when fewer Group Policies are applied. Merge GPOs when possible instead of having multiple GPOs.
Login script execution duration	Indicates the time taken for the login script to execute for this user.	Secs	Compare the value of this measure across users to know for which user the login script took the longest time to

Measurement	Description	Measurement Unit	Interpretation
			execute. If this user complains of slowness, then, you can compare the value of this measure with that of the other 'duration' measures of that user to figure out what could have really caused the slowness.
Start-up client duration	This is the high-level client-side connection start-up metric. It starts at the time of the request (mouse click) and ends when the ICA connection between this user's client device and VDI server has been established.	Secs	In the case of a shared session, this duration will normally be much shorter, as many of the set-up costs associated with the creation of a new connection to the server are not incurred. When any user complains of slowness, you may want to compare the value of this measure with that of the *Session start-up server duration* measure of that user to know whether a client-side issue or a server-side issue is responsible for the slowness. If this comparison reveals that the Start-up client duration of the user is high, it indicates a client-side issue that is causing long start times. In this case therefore, compare the value of the client start-up metrics such as the *Application enumeration client duration, Configuration file download client duration, Credentials obtention client duration, ICA file download duration, Launch page web server duration, Name resolution client duration, Name resolution web server duration, Session look-up client duration, Session creation client duration*, and *Ticket response web server duration* to know what

Measurement	Description	Measurement Unit	Interpretation
			client-side issue is causing the Start-up client duration to be high.
Back-up URL client count	**This measure is relevant when the Citrix receiver is the application launch mechanism**. It records the number of back-up URL retries before a successful launch. Note that this is the only start-up metric that is a measure of attempts, rather than time duration.	Number	If this metric has a value higher than 1, it indicates that the Web Interface server is unavailable and the Citrix receiver (formerly known as Program Neighborhood Agent) is attempting to connect to back-up Web Interface servers to launch the application. A value of 2 means that the main Web Interface server was unavailable, but the Citrix receiver managed to the launch the application successfully using the first back-up server that it tried. A value higher than 2 means that multiple Web Interface servers are unavailable. Probable reasons for the non-availability of the Web Interface servers include (in order of likelihood): • Network issues between the client and the server. So the administrator should make sure that the Web Interface server is on the network and accessible to the clients. • An overloaded Web Interface server that is not responding (or has crashed for another reason). Try to log on to the server and check the Windows Performance Monitor/Task Manager to see how much memory is in use and so on. Also, review the Event Logs to see if Windows logged any serious

Measurement	Description	Measurement Unit	Interpretation
			errors.
Application enumeration client duration	**This measure is relevant when the Citrix receiver is the application launch mechanism**. It measures the time needed by this user's sessions to retrieve the list of applications from the Web Interface service.	Secs	If the *Start-up client duration* measure reports a high value for a user, then compare the value of this measure with that of the other client-side metrics such as the *Configuration file download client duration, Credentials obtention client duration, ICA file download duration, Launch page web server duration, Name resolution client duration, Name resolution web server duration, Session look-up client duration, Session creation client duration*, and *Ticket response web server duration* to know whether/not slowness in application enumeration is the precise reason why it took the user a long time to establish an ICA session with the VDI server.
Configuration file download client duration	**This measure is relevant when the Citrix receiver is the application launch mechanism**. It measures the time this user's sessions took to retrieve the configuration file from the XML server.	Secs	If the *Start-up client duration* measure reports a high value for a user, then compare the value of this measure with that of the other client-side metrics such as *Application enumeration client duration, Credentials obtention client duration, ICA file download duration, Launch page web server duration, Name resolution client duration, Name resolution web server duration, Session look-up client duration, Session creation client duration*, and *Ticket response web server duration* to know whether/not slowness in retrieving the configuration file from the XML server is the precise reason why it took the user a long time an ICA session with

Measurement	Description	Measurement Unit	Interpretation
			the VDI server.
Credentials obtention client duration	**This measure is relevant when the Citrix receiver is the application launch mechanism**. It measures the time required by this user's sessions to obtain the user credentials.	Secs	Note that COCD is only measured when the credentials are entered manually by the user. Because this metric may be artificially inflated if a user fails to provide credentials in a timely manner, it is subtracted from the Start-up client duration measure. However, in the event that the user manually inputs the credentials, and the value of this measure is higher than that of all the other client start-up metrics that this test reports, it is a clear indicator that any connection delay that the user may have experienced is owing to slowness in obtaining user credentials.
ICA file download duration	**This measure is relevant when the Citrix receiver or Web Interface is the application launch mechanism**. This is the time it takes for this user's client to download the ICA file from the web server.	Secs	The overall process here is: a. The user clicks on application icon. b. The user's browser requests the Web Interface launch page. c. The Web Interface launch page receives the request and starts to process the launch, communicating with VDI server and potentially other components such as Secure Ticket Authority (STA). d. The Web Interface generates ICA file data. e. The Web Interface sends the ICA

Measurement	Description	Measurement Unit	Interpretation
			file data back to the user's browser. f. The browser passes ICA file data to the plugin (client). This measure represents the time it takes for the complete process (step 1 to 6). The measure stops counting time when the client receives the ICA file data. The *Launch page web server duration* measure on the other hand, covers the Web server portion of the process (that is, steps 3 and 4). If the ICA file download duration is high, but the Launch page web server duration is normal, it implies that the server-side processing of the launch was successful, but there were communication issues between the client device and the Web server. Often, this results from network trouble between the two machines, so investigate potential network issues first.
Launch page web server duration	This measure is relevant when the Web Interface is the application launch mechanism. It measures the time needed by this user's sessions to process the launch page (launch. aspx) on the Web Interface server.	Secs	If the value of this measure is high, it indicates at a bottleneck on the Web Interface server. Possible causes include: • High load on the Web Interface server. Try to identify the cause of the slow down by checking the Internet Information Services (IIS) logs and monitoring tools, Task Manager, Performance Monitor and

Measurement	Description	Measurement Unit	Interpretation
			so on. • Web Interface is having issues communicating with the other components, such as the VDI server. Check to see if the network connection between Web Interface and VDI is slow or some VDI servers are down or overloaded. If the Web server seems okay, consider reviewing the VDI farm for problems.
Name resolution client duration	This is the time it takes the XML service to resolve the name of a published application to an IP address.	Secs	This metric is collected when a client device directly queries the XML Broker to retrieve published application information stored in IMA (for example, when using the citrix receiver or a Custom ICA Connection). This measure is only gathered for new sessions since session sharing occurs during startup if a session already exists. When this metric is high, it indicates the XML Broker is taking a lot of time to resolve the name of a published application to an IP address. Possible causes include a problem on the client, issues with the XML Broker, such as the XML Broker being overloaded, a problem with the network link between the two, or a problem in IMA. Begin by evaluating traffic on the network and the XML Broker.
Name resolution web server duration	**This measure is relevant when the Citrix receiver or Web Interface is the**	Secs	When this metric is high, there could be an issue with the Web Interface server or the Citrix receiver site

Measurement	Description	Measurement Unit	Interpretation
	application launch mechanism. It is the time it takes the XML service to resolve the name of a published application to a VDI Server address		(formerly known as the Neighborhood Agent site), the XML Service, the network link between the two, or a problem in IMA. Like the *Name resolution client duration* measure, this metric indicates how long it takes the XML service to resolve the name of a published application to a VDI IP address. However, this metric is collected when a Web Interface site is performing this process on behalf of a launch request it has received from either the Citrix receiver (previously known as Program Neighborhood Agent) or from a user clicking a Web Interface page icon. This metric applies to all sessions launched through the Web Interface or the Citrix receiver (formerly, the Program Neighborhood Agent).
Session look-up client duration	Indicates the time this user's sessions take to query every ICA session to host the requested published application.	Secs	The check is performed on the client to determine whether the application launch request can be handled by an existing session. A different method is used depending on whether the session is new or shared.
Session creation client duration	Indicates the new session creation time, from the moment wfica32.exe is launched to the establishment of the connection.	Secs	In the event of slowness, if the Start-up client duration of a user session is found to be higher than the *Session start-up server duration*, you may want to compare the value of this measure with all other client start-up measures to determine whether/not session creation is the process that is slowing down the application launch.
Ticket response web server duration	**This measure is relevant**	Secs	When this metric is high, it can indicate that the Secure Ticket Authority (STA)

Measurement	Description	Measurement Unit	Interpretation
	when the Citrix receiver or Web Interface is the application launch mechanism. This is the time this user's sessions take to get a ticket (if required) from the STA server or XML service.		server or the XML Broker are overloaded.
Reconnect enumeration client duration	**This measure is relevant when the Citrix receiver is the application launch mechanism**. This is the time it takes this user's client to get a list of reconnections.	Secs	Compare the value of this measure with that of other client start-up metrics for a user to know what is the actual cause for the client start-up delay.
Reconnect enumeration web server duration	**This measure is relevant when the Citrix receiver or Web Interface is the application launch mechanism**. This is the time it takes the Web Interface to get the list of reconnections for this user from the XML service.	Secs	Compare the value of this measure with that of other client start-up metrics for a user to know what is the actual cause for the client start-up delay.
Session start-up server duration	This is the high-level server-side connection start-up metric. It includes the time spent on the VDI server to perform the entire start-up operation.	Secs	In the event of an application starting in a shared session, this metric is normally much smaller than when starting a completely new session, which involves potentially high- cost tasks such as profile loading and login script execution. When this metric is high, it indicates that there is a server-side issue increasing session start times. To zero-in on this issue, compare the values of the server start-up metrics such as *Session creation server*

Measurement	Description	Measurement Unit	Interpretation
			duration, Credentials obtention server duration, Program neighborhood credentials obtention server duration, Credentials obtention network server duration, Credentials authentication server duration, Profile load server duration, Login script execution server duration, Drive mapping server duration, and *Printer creation server duration*.
Session creation server duration	Indicates the time spent by the server in creating the session for this user.	Secs	This duration starts when the ICA client connection has been opened and ends when authentication begins. This should not be confused with 'Session start-up server duration'.
Credentials obtention server duration	Indicates the time taken by the server to obtain the credentials of this user.	Secs	This time is only likely to be a significant if manual login is being used and the server-side credentials dialog is displayed (or if a legal notice is displayed before login commences). Because this metric may be artificially inflated if a user fails to provide credentials in a timely manner, it is not included in the Credentials obtention server duration. However, in the event that the user manually inputs the credentials, and the value of this measure is higher than that of all the other client start-up metrics that this test reports, it is a clear indicator that any connection delay that the user may have experienced is owing to slowness in obtaining user credentials.
Credentials obtention network server duration	Indicates the time spent by the server performing network operations to	Secs	This only applies to a Security Support Provider Interface login (a form of pass-through authentication where the

Measurement	Description	Measurement Unit	Interpretation
	obtain credentials for this user.		client device is a member of the same domain as the server and Kerberos tickets are passed in place of manually entered credentials).
Program neighborhood credentials obtention server duration	This only applies to a Security Support Provider Interface login (a form of pass-through authentication where the client device is a member of the same domain as the server and Kerberos tickets are passed in place of manually entered credentials).	Secs	As in the case of the Credentials obtention server duration metric, because this metric may be artificially inflated if a user fails to provide credentials in a timely manner, it is not included in the Credentials obtention server duration.
Credentials authentication server duration	Indicates the time spent by the server when authenticating the user's credentials against the authentication provider, which may be Kerberos, Active Directory or a Security Support Provider Interface (SSPI).	Secs	Where server-side issues are causing user experience to deteriorate, you can compare the value of this measure with that of all the other server start-up metrics that this test reports - i.e., Session creation server duration, Credentials obtention server duration, Program neighborhood credentials obtention server duration, Credentials obtention network server duration, Profile load server duration, Login script execution server duration, Drive mapping server duration, and Printer creation server duration - to know what is the root-cause of delays in server start-up.
Profile load server duration	Indicates the time required by the server to load this user's profile.	Secs	If this metric is high, consider your Terminal Services profile configuration. Citrix Consulting has found that when customers have logon times greater than 20 seconds, in most cases, this can be attributed to poor profile and policy design. Roaming profile size and

Measurement	Description	Measurement Unit	Interpretation
			location contribute to slow session starts. When a user logs onto a session where Terminal Services roaming profiles and home folders are enabled, the roaming profile contents and access to that folder are mapped during logon, which takes additional resources. In some cases, this can consume significant amounts of the CPU usage. Consider using the Terminal Services home folders with redirected personal folders to mitigate this problem. In general, consider using Citrix Profile management to manage user profiles in Citrix environments. This tool also provides logging capabilities to help isolate profile issues. If you are using Citrix profile management and have slow logon times, check to see if your antivirus software is blocking the Citrix profile management tool.
Login script execution server duration	Indicates the time needed by the server to run this user's login script(s).	Secs	If the value of this measure is abnormally high for any user, consider if you can streamline this user or group's login scripts. Also, consider if you can optimize any application compatibility scripts or use environment variables instead.
Drive mapping server duration	Indicates the time needed for the server to map this user's client drives, devices and ports.	Secs	Make sure that, when possible, your base policies include settings to disable unused virtual channels, such as audio or COM port mapping, to optimize the ICA protocol and improve overall session performance.
Printer creation	Indicates the time required	Secs	If the configuration is set such that

Measurement	Description	Measurement Unit	Interpretation
server duration	for the server to synchronously map this user's client printers.		printer creation is performed asynchronously, no value is recorded for this measure as it is does not impact completion of the session start-up. On the other hand, if excessive time is spent mapping printers, it is often the result of the printer autocreation policy settings. The number of printers added locally on the users' client devices and your printing configuration can directly affect your session start times. When a session starts, VDI has to create every locally mapped printer on the client device. Consider reconfiguring your printing policies to reduce the number of printers that get created - especially if users have a lot of local printers.

Chapter 4: Monitoring VDI-in-a-Box On XenServer

To monitor the Citrix XenServer and the vdimanager operating on it, use the VDI in a Box / XenServer monitoring model.

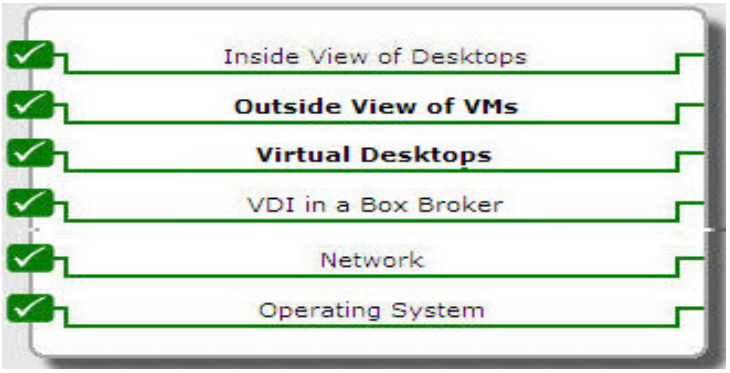

Figure 4.1: Layer model of VDI in a Box/XenServer

Each layer of this model is mapped to tests that report a variety of statistics related to the Citrix XenServer and the VDI in a box appliance running on it. While the two layers at the bottom of the layer model focus on the health of the operating system and the network of the Xen host, the top two layers perform 'In-N-Out' monitoring of the virtual desktops operating on the target XenServer. To ascertain the status of the vdimanager and its operations on the other hand, you need to use the Virtual Desktops and VDI in a Box Broker layers.

To enable the eG remote agent to pull out all these performance statistics, you first need to make sure that the following pre-requisites are in place:

For monitoring the vdimanager:

- The eG agent should be able to connect to the vdimanager via SSH and pull out the metrics. For this purpose, the default SSH port, *22*, should be opened on the vdimanager. If your environment has been configured with a different SSH port, then make sure that port is open.
- The eG agent should be able to login to the vdimanager appliance for monitoring and metrics collection. To enable this, you need to configure all tests that the eG agent executes on the vdimanager with the credentials of a user with login rights. By default, the appliance supports a root user named root and a user named **kvm**. If you prefer not to expose the credentials of the root user, then you can configure the tests with the credentials of the other user **kvm** for this purpose. By default, the user root takes the password root, and the user **kvm** takes the password **kaviza123**.

For 'In-N-Out' monitoring of the Citrix XenServer and its VMs:

- Make sure that the pre-requisites pertaining to 'agentless' monitoring of XenServers (detailed in the *Monitoring XenServers* document) are fulfilled.

For details on the **Virtual Desktops** and **VDI In a Box Broker** layers refer to the **Managing the VDI-in-a-Box on VMware vSphere** chapter. For all the other layers, refer to the *Monitoring XenServers* document.

Chapter 5: Monitoring VDI-in-a-Box On Hyper-V

To monitor Microsoft Hyper-V and the vdimanager operating on it, use the VDI-in-a-Box / Hyper-V monitoring model.

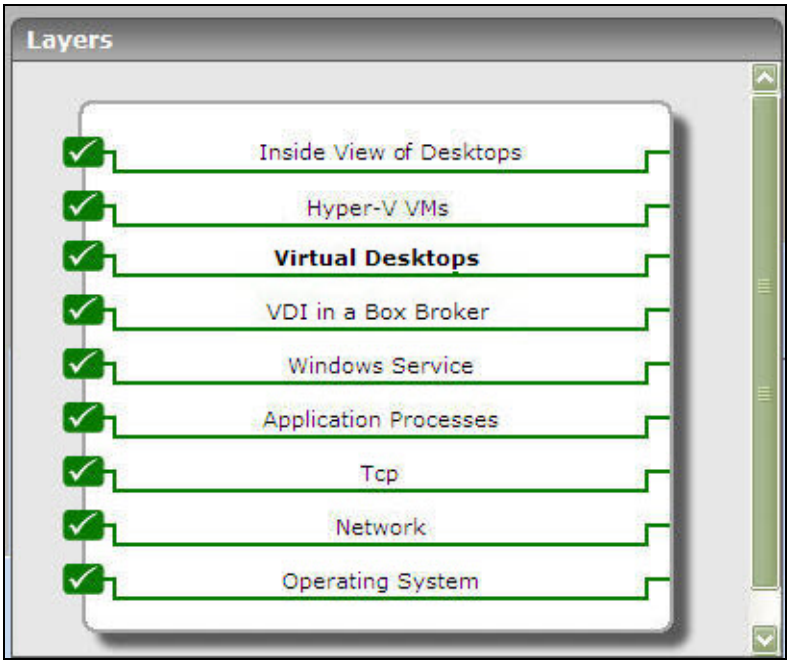

Figure 5.1: Layer model of VDI in a Box/Hyper-V

Each layer of this model is mapped to tests that report a variety of statistics related to the Microsoft Hyper-V and the VDI in a box appliance running on it. While the five layers at the bottom of the layer model focus on the health of the operating system and the network of the Hyper-V host, the top two layers perform 'In-N-Out' monitoring of the virtual desktops operating on the target Hyper-V server. To ascertain the status of the vdimanager and its operations on the other hand, you need to use the **Virtual Desktops** and **VDI in a Box Broker** layers.

To enable the eG remote agent to pull out all these performance statistics, you first need to make sure that the following pre-requisites are in place:

For monitoring the vdimanager:

- The eG agent should be able to connect to the vdimanager via SSH and pull out the metrics. For this purpose, the default SSH port, 22, should be opened on the vdimanager. If your environment has been configured with a different SSH port, then make sure that port is open.

- The eG agent should be able to login to the vdimanager appliance for monitoring and metrics collection. To enable this, you need to configure all tests that the eG agent executes on the vdimanager with the credentials of a user with login rights. By default, the appliance supports a root user named **root** and a user named **kvm**. If you prefer not to expose the credentials of the root user, then you can configure the tests with the credentials of the other user **kvm** for this purpose. By default, the user **root** takes the password **root**, and the user **kvm** takes the password **kaviza123**.

For 'In-N-Out' monitoring of the Microsoft Hyper-V and its VMs:

- Make sure that the pre-requisites pertaining to 'agentless' monitoring of the Hyper-V server (detailed in Section 1.3 of the *Monitoring Microsoft Hyper-V* document) are fulfilled.

For details on the **Virtual Desktops** and **VDI In a Box Broker** layers refer to **Managing the VDI-in-a-Box on VMware vSphere**. For all the other layers, refer to the *Monitoring Microsoft Hyper-V* document.

About eG Innovations

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

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

To learn more visit www.eginnovations.com.

Contact Us

For support queries, email support@eginnovations.com.

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