Monitoring Microsoft RDS Server

eG Innovations Product Documentation

www.eginnovations.com

Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

The Microsoft RDS Server is a server program that provides the graphical user interface (GUI) of the Windows desktop to user terminals that don't have this capability themselves. The latter include the relatively low-cost NetPC or "thin client" that some companies are purchasing as alternatives to the autonomous and more expensive PC with its own operating system and applications.

Typically, Microsoft RDS server environments involve multiple tiers of software. Domain servers in the target infrastructure handle authentication of users. Authenticated requests are passed to the Microsoft RDS servers that host a number of applications. In turn, the applications may use backend databases, printers, etc., for different functionalities. Owing to the multi-tier nature of Microsoft RDS server environments, a slow-down in one tier (e.g., the authentication server) can cause a slow-down of the entire service. When a slow-down occurs, an administrator of the server farm has to quickly determine what the source of the problem could be - i.e., Is it the network? Or the authentication server? Or the Microsoft RDS server? Or the backend database? Or the application? Accurate, fast diagnosis of problems helps reduce downtime and improve customer satisfaction. The eG Enterprise Suite helps administrator in this regards!

Chapter 2: Administering eG Manager to work with Microsoft RDS

To do the above, do the following:

1. Log into the eG administrative interface.
2. If a Microsoft RDS server is already discovered, then directly proceed towards managing it using the **COMPONENTS - MANAGE/UNMANAGE** page (Infrastructure -> Components -> Manage/Unmanage). However, if it is yet to be discovered, then run discovery (Infrastructure -> Components -> Discovery) to get it discovered or add the Microsoft RDS server manually using the **COMPONENTS** page (Infrastructure -> Components -> Add/Modify). Remember that components manually added are managed automatically. Discovered components, however, are managed using the **COMPONENTS - MANAGE/UNMANAGE** page. Figure 2.1 and Figure 2.2 clearly illustrate the process of managing a WINS server.

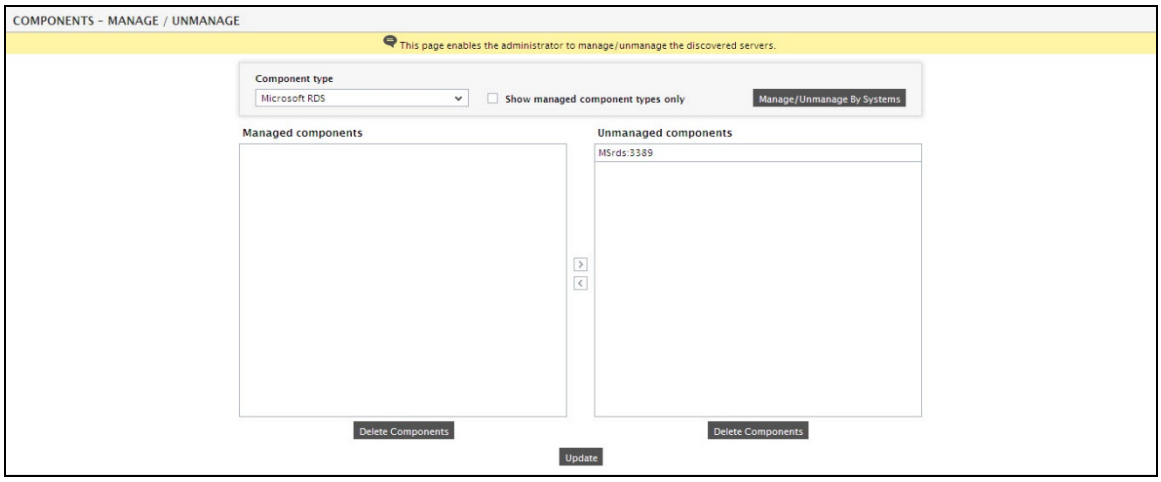

Figure 2.1: Selecting the Microsoft RDS server to manage

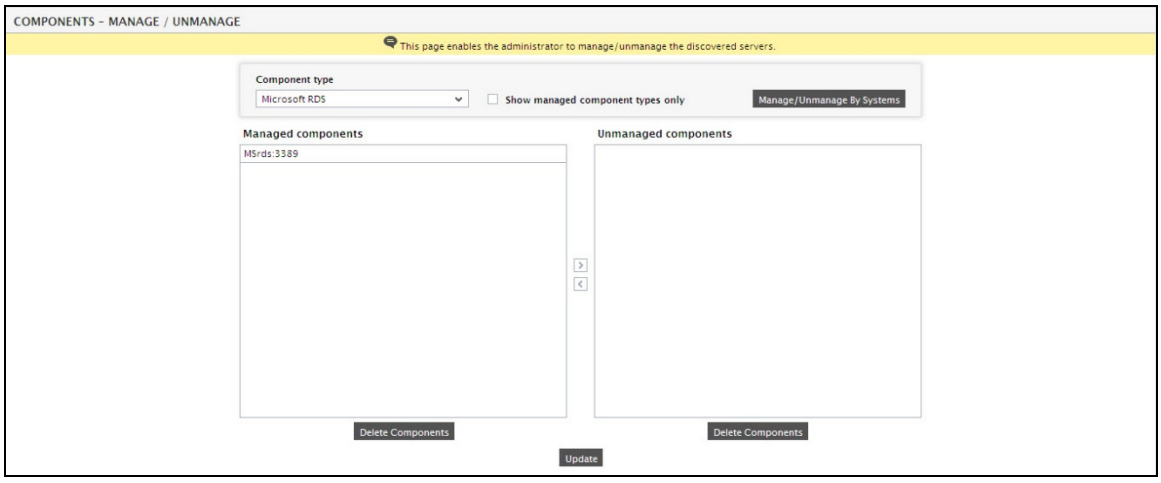

Figure 2.2: Managing the Microsoft RDS server

3. Next, try to sign out of the eG administrative interface. Doing so will result in the List of unconfigured tests listing the unconfigured tests of the Microsoft RDS server (see Figure 2.3).

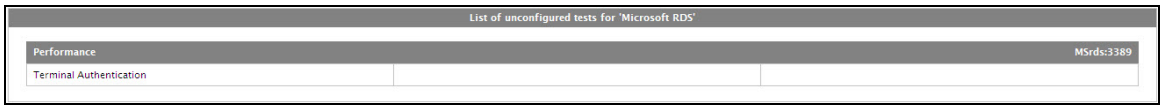

Figure 2.3: Unconfigured tests of the Microsoft RDS server

4. Now, click on the **Terminal Authentication** test to configure it. To know how to configure the test, refer to Section **3.3.3**.

Chapter 3: Monitoring Microsoft RDS Servers

The Microsoft RDS Server is a server program that provides the graphical user interface (GUI) of the Windows desktop to user terminals that don't have this capability themselves. The latter include the relatively low-cost NetPC or "thin client" that some companies are purchasing as alternatives to the autonomous and more expensive PC with its own operating system and applications.

Typically, Microsoft RDS server environments involve multiple tiers of software. Domain servers in the target infrastructure handle authentication of users. Authenticated requests are passed to the Microsoft RDS servers that host a number of applications. In turn, the applications may use backend databases, printers, etc., for different functionalities. Owing to the multi-tier nature of Microsoft RDS server environments, a slow-down in one tier (e.g., the authentication server) can cause a slow-down of the entire service. When a slow-down occurs, an administrator of the server farm has to quickly determine what the source of the problem could be - i.e., Is it the network? Or the authentication server? Or the Microsoft RDS server? Or the backend database? Or the application? Accurate, fast diagnosis of problems helps reduce downtime and improve customer satisfaction.

The eG Enterprise suite offers 100% web-based monitoring of Microsoft RDS server farms. The suite includes an extensive, pre-defined, customized Microsoft RDS model for this server (see Figure 3.1), which defines the key performance metrics that need to be tracked to determine the service level achieved by the server/server farm.

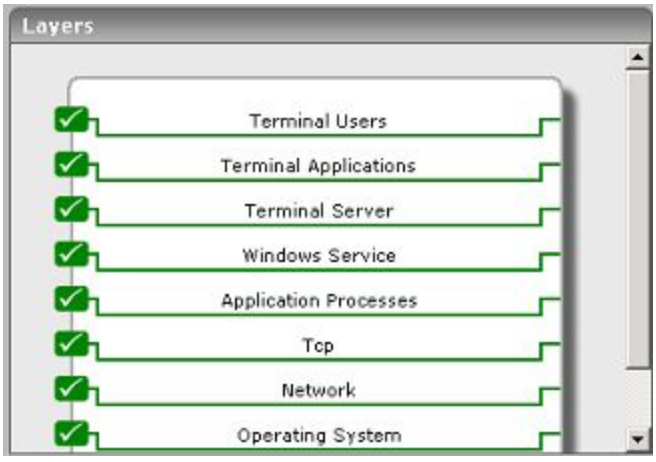

Figure 3.1: Layer model of a Microsoft RDS server

Using the metrics reported by each of the layers depicted by Figure 3.1, administrators can find answers to persistent performance-related queries discussed hereunder:

Microsoft RDS server Monitoring	• Are the Microsoft RDS servers available to service user requests? • Are there sporadic disconnects from the Microsoft RDS server? • At what times do peak usage of the servers happen and is the server capacity adequate? • Is the user load being balanced across all the servers? • Is the data store available?
User Monitoring	• What is the average response time that a user sees when connecting to a Microsoft RDS server? • How many users are logged in to each server in the Microsoft RDS server farm? What is the resource usage (CPU and memory) for each user? • What is the I/O activity generated by every user? • How much network bandwidth is consumed by every user? • Are too many page faults occurring in the processes executed on a server? • If so, what are those processes, and who are the users executing them? • Which user is using a lot of handles?
Operating System Monitoring	• What is the average CPU and memory usage on all the servers in the farm? • Is any unusual memory scanning/paging activity happening on the systems? • Are the critical Microsoft RDS server processes up? • What is their resource consumption?
Hosted Application Monitoring	• What are the applications hosted on a Microsoft RDS server? • Who is using each application? • What is the resource usage for each published application?
Infrastructure Services Monitoring	• Are the backend databases working? • What is the resource usage of the databases?

	• Are users able to login to the server farm? How long is the login process taking? • What is the usage of the Microsoft Windows Domain Controller?

Since the 4 layers at the bottom of Figure 3.1 have already been discussed in the *Monitoring Unix and Windows Servers* document, the sections to come will discuss the top 4 layers only.

3.1 The Operating System Layer

This layer measures the health of the Windows OS on which Microsoft RDS operates. Typically, all the tests mapped to the **Operating System** layer of any managed Windows server in the environment will be mapped to the **Operating System** layer of a Microsoft RDS component as well. To know the details of these tests, refer to the *Monitoring Unix and Windows Servers* document.

The only additional test that runs at the Operating System layer of the Microsoft RDS component is the **Grid GPUs** test. This test has been discussed below.

3.1.1 Grid GPUs Test

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

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

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

Virtual application delivery with XenApp/RDS and NVIDIA GRID™ offloads graphics processing from the CPU to the GPU, allowing the data center manager to deliver to all user types for the first time.

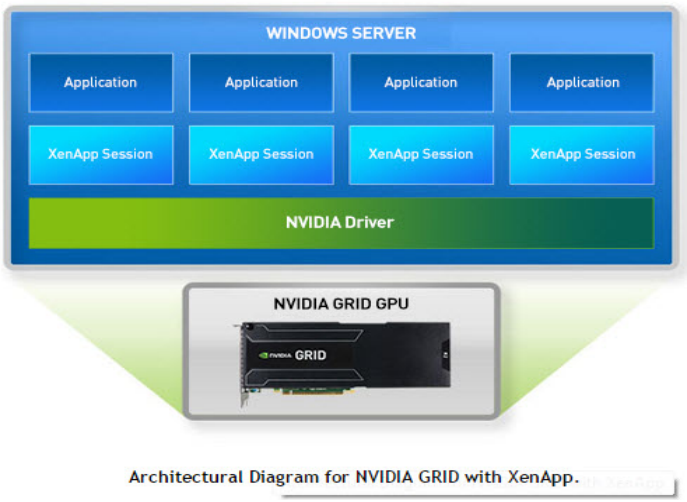

Figure 3.2: Architectural diagram for NVIDIA GRID with XenApp

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

- The host is sized with sufficient GPU resources;
- The GPUs are configured with enough graphics memory;

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

Using this test, administrators can identify the physical GPUs on the NVIDIA GRID card used by the host. For each physical GPU, administrators can determine how actively memory on that GPU is utilized, thus revealing the GPU on which memory is used consistently. In addition, the test also indicates how busy each GPU is, and in the process pinpoints those physical GPUs that are being over-utilized by the virtual applications on the host. The adequacy of the physical GPU resources is thus revealed. Moreover, the power consumption and temperature of each GPU of the host is also monitored and its current temperature and power usage can be ascertained; administrators are thus alerted to abnormal power usage of the GPU and unexpected fluctuations in its temperature. The power limit set and the clock frequencies configured are also revealed, so that administrators can figure out whether the GPU is rightly configured for optimal processing or is any fine-tuning required.

Note:

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

For this test to run and report metrics, the NVWMI should be installed on the Citrix XenApp server. To know how, refer to the Configuring the eG Agent to Monitor NVIDIA Graphics Processing Units (GPUs) section of the Monitoring Citrix XenServers document.

Target of the test : A Citrix XenApp server / Microsoft RDS server

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for each GRID physical GPU assigned to the host 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 the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server.
NVIDIA Path	Specify the full path to the install directory of the NVIDIA. By default, the NVIDIA will be installed in the *C:/Progra~1/NVIDIA~1/NVSMIdirectory*. If the NVIDIA indeed resides in its default location, set the nvidia path to none. On the other hand, if the NVIDIA has been installed in a different location, provide the full path to that location against nvidia path.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
GPU memory utilization	Indicates the proportion of time over the past sample period during which global (device) memory was being read or written on this GPU.	Percent	A value close to 100% is a cause for concern as it indicates that graphics memory on a GPU is almost always in use. In a XenApp/RDS environment, this could be because one/more sessions to XenApp are consistently

Measurement	Description	Measurement Unit	Interpretation
			accessing graphic- intensive applications. If the value of this measure is high almost all the time for most of the GPUs, it could mean that the host is not sized with adequate graphics memory.
Used frame buffer memory	Indicates the amount of frame buffer memory on-board this GPU that is being used by the host.	MB	Frame buffer memory refers to the memory used to hold pixel properties such as color, alpha, depth, stencil, mask, etc. Properties like the screen resolution, color level, and refresh speed of the frame buffer can impact graphics performance. Also, if Error-correcting code (ECC) is enabled on a host, the available frame buffer memory may be decreased by several percent. This is because, ECC uses up memory to detect and correct the most common kinds of internal data corruption. Moreover, the driver may also reserve a small amount of memory for internal use, even without active work on the GPU; this too may impact frame buffer memory. For optimal graphics performance therefore, adequate frame buffer memory should be allocated to the host.
Free frame buffer memory	Indicates the amount of frame buffer memory	MB	

Measurement	Description	Measurement Unit	Interpretation
	on-board this GPU that is yet to be used by the host.		
Frame buffer memory utilization	Indicates the percentage of frame buffer memory on-board this GPU that is being utilized by the host.	Percent	A value close to 100% is indicative of excessive frame buffer memory usage. Properties like the screen resolution, color level, and refresh speed of the frame buffer can impact graphics performance. Also, if Error-correcting code (ECC) is enabled on a host, the available frame buffer memory may be decreased by several percent. This is because, ECC uses up memory to detect and correct the most common kinds of internal data corruption. Moreover, the driver may also reserve a small amount of memory for internal use, even without active work on the GPU; this too may impact frame buffer memory. For optimal graphics performance therefore, adequate frame buffer memory should be allocated to the host.
GPU compute utilization	Indicates the proportion of time over the past sample period during which one or more kernels was executing on this GPU.	Percent	A value close to 100% indicates that the GPU is busy processing graphic requests almost all the time. In a XenApp/RDS environment, this could be because one/more sessions to XenApp are consistently accessing graphic- intensive

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

Measurement	Description	Measurement Unit	Interpretation
	GPU that is currently being used by the host.		memory should be available to the host.
Free BAR1 memory	Indicates the amount of BAR1 memory of this GPU that is still to be used by the host.	MB	
BAR1 memory utilization	Indicates the percentage of the total BAR1 memory on this GPU that is currently being utilized by the host.	Percent	A value close to 100% is indicative of excessive BAR1 memory usage by the host. For best graphics performance, sufficient BAR1 memory resources should be available to the host.
Power management	Indicates whether/not power management is enabled for this GPU.		Many NVIDIA graphics cards support multiple performance levels so that the server can save power when full graphics performance is not required. The default Power Management Mode of the graphics card is Adaptive. In this mode, the graphics card monitors GPU usage and seamlessly switches between modes based on the performance demands of the application. This allows the GPU to always use the minimum amount of power required to run a given application. This mode is recommended by NVIDIA for best overall balance of power and performance. If the power management mode is set to Adaptive, the value of this measure will be Supported.

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>Alternatively, you can set the Power Management Mode to Maximum Performance. This mode allows users to maintain the card at its maximum performance level when 3D applications are running regardless of GPU usage. If the power management mode of a GPU is Maximum Performance, then the value of this measure will be Maximum.

The numeric values that correspond to these measure values are discussed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Supported</td><td>1</td></tr><tr><td>Maximum</td><td>0</td></tr></table>
Note:

By default, this measure will report the Measure Values listed in the table above to indicate the power management status. In the graph of this measure however, the same is represented using the numeric equivalents only.</td></tr>
<tr><td>Power limit</td><td>Indicates the power limit configured for this GPU.</td><td>Watts</td><td>This measure will report a value only if the value of the ‘Power management’ measure is ‘Supported’.

The power limit setting controls how much voltage a GPU can use when under load. Its not advisable to set</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			the power limit at its maximum – i.e., the value of this measure should not be the same as the value of the Max power limit measure - as it can cause the GPU to behave strangely under duress.
Default power limit	Indicates the default power management algorithm's power ceiling for this GPU.	Watts	**This measure will report a value only if the value of the 'Power management' measure is 'Supported'.**
Enforced power limit	Indicates the power management algorithm's power ceiling for this GPU.	Watts	**This measure will report a value only if the value of the 'Power management' measure is 'Supported'.** The total board power draw is manipulated by the power management algorithm such that it stays under the value reported by this measure.
Min power limit	The minimum value that the power limit of this GPU can be set to.	Watts	This measure will report a value only if the value of the 'Power management' measure is 'Supported'.
Max power limit	The maximum value that the power limit of this GPU can be set to.	Watts	**This measure will report a value only if the value of the 'Power management' measure is 'Supported'.** If the value of this measure is the same as that of the Power limit measure, then the GPU may behave strangely.

Measurement	Description	Measurement Unit	Interpretation
Graphics clock	Indicates the current frequency of the graphics clock of this GPU.	MHz	GPU has many more cores than your average CPU but these cores are much simpler and much smaller so that many more actually fit on a small piece of silicon. These smaller, simpler cores go by different names depending upon the tasks they perform. Stream processors are the cores that perform a single thread at a slow rate. But since GPUs contain numerous stream processors, they make overall computation high. The streaming multiprocessor clock refers to how fast the stream processors run. The Graphics clock is the speed at which the GPU operates. The memory clock is how fast the memory on the card runs. By correlating the frequencies of these clocks (i.e., the value of these measures) with the memory usage, power usage, and overall performance of the GPU, you can figure out if overclocking is required or not. Overclocking is the process of forcing a GPU core/memory to run faster than its manufactured frequency. Overclocking can have both positive and negative effects on GPU performance. For instance, memory overclocking helps on cards with low memory bandwidth, and with games with a lot of post-

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

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

Measurement	Description	Measurement Unit	Interpretation
			should be 0.
Volatile double bit errors	Indicates the total number of volatile double bit errors in this GPU.	Number	Volatile error counters track the number of errors detected since the last driver load. Double bit errors are detected but not corrected. Ideally, the value of this measure should be 0.
Aggregate single bit errors	Indicates the total number of aggregate single bit errors in this GPU.	Number	Aggregate error counts persist indefinitely and thus act as a lifetime counter. Single bit ECC errors are automatically corrected by the hardware and do not result in data corruption. Ideally, the value of this measure should be 0.
Aggregate double bit errors	Indicates the total number of aggregate double bit errors in this GPU.	Number	Aggregate error counts persist indefinitely and thus act as a lifetime counter. Double bit errors are detected but not corrected. Ideally, the value of this measure should be 0.

3.2 The Windows Service Layer

This layer represents the different services of the corresponding Windows components in the environment. An eG agent uses **Windows Services** test to track the health of this layer. In addition, the layer also periodically monitors the application, security, and system-related events that occur on the target Windows host. Since most of the tests of this layer have already been dealt in the Monitoirng Unix and Windows servers document, let us now discuss the tests that are exclusive for the Microsoft RDS Servers alone.

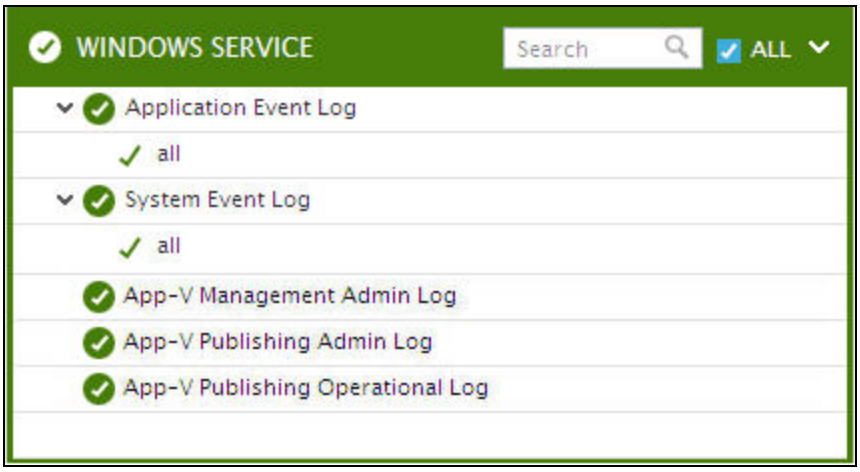

Figure 3.3: The tests mapped to the Windows Service layer

3.2.1 App-V Client Admin Log Test

This test reports the statistical information about the admin events generated by the target system.

Note:

This test will report metrics only when the App-V Client is installed on the Microsoft RDS Server.

Target of the test : An App-V Client on the target Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the App-V Client that is to be monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. By default, this is *8080*.
Logtype	Refers to the type of event logs to be monitored. The default value is **Microsoft-AppV-Client/Admin**.
Policy Based Filter	Using this page, administrators can configure the event sources, event IDs, and event descriptions to be monitored by this test. In order to enable administrators to easily and accurately provide this specification, this page provides the following options:

Parameters	Description
	• Manually specify the event sources, IDs, and descriptions in the Filter text area, or, • Select a specification from the predefined filter policies listed in the Filter box For explicit, manual specification of the filter conditions, select the **No** option against the Policy Based Filter field. This is the default selection. To choose from the list of pre-configured filter policies, or to create a new filter policy and then associate the same with the test, select the **Yes** option against the Policy Based Filter field.
Filter	If the Policy Based Filter flag is set to **No**, then a Filter text area will appear, wherein you will have to specify the event sources, event IDs, and event descriptions to be monitored. This specification should be of the following format: *{Displayname}:{event_ sources_ to_ be_ included}:{event_ sources_ to_ be_ excluded}: {event_ IDs_ to_ be_ included}: {event_ IDs_ to_ be_ excluded}: {event_ descriptions_ to_ be_ included}:{event_ descriptions_ to_ be_ excluded}*. For example, assume that the Filter text area takes the value, *OS_ events:all:Browse,Print:all:none:all:none*. Here: • *OS_events* is the display name that will appear as a descriptor of the test in the monitor UI; • *all* indicates that all the event sources need to be considered while monitoring. To monitor specific event sources, provide the source names as a comma-separated list. To ensure that none of the event sources are monitored, specify *none*. • Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, *Browse* and *Print* have been excluded from monitoring. Alternatively, you can use *all* to indicate that all the event sources have to be excluded from monitoring, or *none* to denote that none of the event sources need be excluded. • In the same manner, you can provide a comma-separated list of event IDs that require monitoring. The *all* in our example represents that all the event IDs need to be considered while monitoring. • Similarly, the *none* (following *all* in our example) is indicative of the fact that none of the event IDs need to be excluded from monitoring. On the other hand, if you want to instruct the eG Enterprise system to ignore a few event IDs during monitoring, then

Parameters	Description
	provide the IDs as a comma-separated list. Likewise, specifying *all* makes sure that all the event IDs are excluded from monitoring. • The *all* which follows implies that all events, regardless of description, need to be included for monitoring. To exclude all events, use *none*. On the other hand, if you provide a comma-separated list of event descriptions, then the events with the specified descriptions will alone be monitored. Event descriptions can be of any of the following forms - *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. • In the same way, you can also provide a comma-separated list of event descriptions to be excluded from monitoring. Here again, the specification can be of any of the following forms: *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. In our example however, none is specified, indicating that no event descriptions are to be excluded from monitoring. If you use *all* instead, it would mean that all event descriptions are to be excluded from monitoring. By default, the Filter parameter contains the value: all. Multiple filters are to be separated by semi-colons (;). **Note:** The event sources and event IDs specified here should be exactly the same as that which appears in the Event Viewer window. On the other hand, if the Policy Based Filter flag is set to **Yes**, then a Filter list box will appear, displaying the filter policies that pre-exist in the eG Enterprise system. A filter policy typically comprises of a specific set of event sources, event IDs, and event descriptions to be monitored. This specification is built into the policy in the following format: *{Policyname}: {event_ sources_ to_ be_ included}: {event_ sources_ to_ be_ excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}* To monitor a specific combination of event sources, event IDs, and event

Parameters	Description
	descriptions, you can choose the corresponding filter policy from the Filter list box. Multiple filter policies can be so selected. Alternatively, you can modify any of the existing policies to suit your needs, or create a new filter policy. To facilitate this, a  icon appears near the Filter list box, once the **Yes** option is chosen against Policy Based Filter. Clicking on the icon leads you to a page where you can modify the existing policies or create a new one. The changed policy or the new policy can then be associated with the test by selecting the policy name from the Filter list box in this page.
UseWMI	The eG agent can either use **WMI** to extract event log statistics or directly parse the event logs using **event log API**s. If the UseWMI flag is **Yes**, then **WMI** is used. If not, the **event log API**s are used. This option is provided because on some Windows NT/2000 systems (especially ones with service pack 3 or lower), the use of WMI access to event logs can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the UseWMI parameter value to **No**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the UseWMI flag should always be set to 'Yes'.**
Stateless Alerts	Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the **EventLog** test. When this test captures an error event for the very first time, the eG manager will send out a **Critical** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **Critical**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting** capability. To enable this capability for this test, set the Stateless Alerts flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.
DDforinformation	eG Enterprise also provides you with options to restrict the amount of storage

Parameters	Description
	required for event log tests. Towards this end, the DDforinformation and DDforwarning flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the DDforinformation flag to **No**.
DDforwarning	To ensure that the test does not generate and store detailed measures for warning events, set the DDforwarning flag to **No**.
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 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
Information messages	Indicates the number of App- V Client admin information events	Number	A change in the value of this measure may indicate infrequent but successful operations

Measurement	Description	Measurement Unit	Interpretation
	generated when the test was last executed.		performed by one or more applications. Please check the App-V Client admin logs in the Event Log Viewer for more details.
Warnings	Indicates the number of App-V Client admin warnings that were generated when the test was last executed.	Number	A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more applications. Please check the App-V Client admin logs in the Event Log Viewer for more details.
Error messages	Indicates the number of App-V Client admin error events that were generated during the last measurement period.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more applications. Please check the App-V Client admin logs in the Event Log Viewer for more details.
Critical messages	Indicates the number of App-V Client admin critical error events that were generated when the test was last executed.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of fatal/irrepairable problems in one or

Measurement	Description	Measurement Unit	Interpretation
			more applications. Please check the App- V Client admin logs in the Event Log Viewer for more details.
Verbose messages	Indicates the number of App- V Client admin verbose events that were generated when the test was last executed.	Number	The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period. Please check the App- V Client admin logs in the Event Log Viewer for more details.

3.2.2 App-V Client Operational Log Test

This test reports the statistical information about the operation events generated by the target system.

Note:

This test will report metrics only when the App-V Client is installed on the Microsoft RDS Server.

Target of the test : An App-V Client on the target Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the App-V Client that is to be monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. By default, this is *8080*.
Logtype	Refers to the type of event logs to be monitored. The default value is **Microsoft-AppV-Client/Operational**.
Policy Based Filter	Using this page, administrators can configure the event sources, event IDs, and

Parameters	Description
	event descriptions to be monitored by this test. In order to enable administrators to easily and accurately provide this specification, this page provides the following options: • Manually specify the event sources, IDs, and descriptions in the Filter text area, or, • Select a specification from the predefined filter policies listed in the Filter box For explicit, manual specification of the filter conditions, select the **No** option against the Policy Based Filter field. This is the default selection. To choose from the list of pre-configured filter policies, or to create a new filter policy and then associate the same with the test, select the **Yes** option against the Policy Based Filter field.
Filter	If the Policy Based Filter flag is set to **No**, then a Filter text area will appear, wherein you will have to specify the event sources, event IDs, and event descriptions to be monitored. This specification should be of the following format: *{Displayname}:{event_ sources_ to_ be_ included}:{event_ sources_ to_ be_ excluded}: {event_ IDs_ to_ be_ included}: {event_ IDs_ to_ be_ excluded}: {event_ descriptions_ to_ be_ included}:{event_ descriptions_ to_ be_ excluded}*. For example, assume that the Filter text area takes the value, *OS_ events:all:Browse,Print:all:none:all:none*. Here: • *OS_events* is the display name that will appear as a descriptor of the test in the monitor UI; • *all* indicates that all the event sources need to be considered while monitoring. To monitor specific event sources, provide the source names as a comma-separated list. To ensure that none of the event sources are monitored, specify *none*. • Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, *Browse* and *Print* have been excluded from monitoring. Alternatively, you can use *all* to indicate that all the event sources have to be excluded from monitoring, or *none* to denote that none of the event sources need be excluded. • In the same manner, you can provide a comma-separated list of event IDs that require monitoring. The *all* in our example represents that all the event IDs need to be considered while monitoring.

Parameters	Description
	• Similarly, the *none* (following *all* in our example) is indicative of the fact that none of the event IDs need to be excluded from monitoring. On the other hand, if you want to instruct the eG Enterprise system to ignore a few event IDs during monitoring, then provide the IDs as a comma-separated list. Likewise, specifying *all* makes sure that all the event IDs are excluded from monitoring. • The *all* which follows implies that all events, regardless of description, need to be included for monitoring. To exclude all events, use *none*. On the other hand, if you provide a comma-separated list of event descriptions, then the events with the specified descriptions will alone be monitored. Event descriptions can be of any of the following forms - *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. • In the same way, you can also provide a comma-separated list of event descriptions to be excluded from monitoring. Here again, the specification can be of any of the following forms: *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. In our example however, none is specified, indicating that no event descriptions are to be excluded from monitoring. If you use *all* instead, it would mean that all event descriptions are to be excluded from monitoring. By default, the Filter parameter contains the value: all. Multiple filters are to be separated by semi-colons (;). **Note:** The event sources and event IDs specified here should be exactly the same as that which appears in the Event Viewer window. On the other hand, if the Policy Based Filter flag is set to **Yes**, then a Filter list box will appear, displaying the filter policies that pre-exist in the eG Enterprise system. A filter policy typically comprises of a specific set of event sources, event IDs, and event descriptions to be monitored. This specification is built into the policy in the following format:

Parameters	Description
	{Policyname}: {event_ sources_ to_ be_ included}: {event_ sources_ to_ be_ excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_ descriptions_to_be_included}:{event_descriptions_to_be_excluded} To monitor a specific combination of event sources, event IDs, and event descriptions, you can choose the corresponding filter policy from the Filter list box. Multiple filter policies can be so selected. Alternatively, you can modify any of the existing policies to suit your needs, or create a new filter policy. To facilitate this, a icon appears near the Filter list box, once the **Yes** option is chosen against Policy Based Filter. Clicking on the icon leads you to a page where you can modify the existing policies or create a new one. The changed policy or the new policy can then be associated with the test by selecting the policy name from the Filter list box in this page.
UseWMI	The eG agent can either use **WMI** to extract event log statistics or directly parse the event logs using **event log API**s. If the UseWMI flag is **Yes**, then **WMI** is used. If not, the **event log API**s are used. This option is provided because on some Windows NT/2000 systems (especially ones with service pack 3 or lower), the use of WMI access to event logs can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the UseWMI parameter value to **No**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the UseWMI flag should always be set to 'Yes'.**
Stateless Alerts	Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the **EventLog** test. When this test captures an error event for the very first time, the eG manager will send out a **Critical** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **Critical**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting** capability. To enable this capability for this test, set the Stateless Alerts flag to

Parameters	Description
	Yes. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.
DDforinformation	eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the DDforinformation and DDforwarning flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the DDforinformation flag to **No**.
DDforwarning	To ensure that the test does not generate and store detailed measures for warning events, set the DDforwarning flag to **No**.
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 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
Information messages	Indicates the number of App-V Client operational information events generated when the test was last executed.	Number	A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications. Please check the App-V Client Operational logs in the Event Log Viewer for more details.
Warnings	Indicates the number of App-V Client operational warnings that were generated when the test was last executed.	Number	A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more applications. Please check the App-V Client Operational logs in the Event Log Viewer for more details.
Error messages	Indicates the number of App-V Client operational error events that were generated during the last measurement period.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more applications. Please check the App-V Client Operational logs in the Event Log Viewer for more details.
Critical messages	Indicates the number of App-V Client operational critical error events that were generated when	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential

Measurement	Description	Measurement Unit	Interpretation
	the test was last executed.		problems. An increasing trend or high value indicates the existence of fatal/irrepairable problems in one or more applications. Please check the App- V Client Operational logs in the Event Log Viewer for more details.
Verbose messages	Indicates the number of App-V Client operational verbose events that were generated when the test was last executed.	Number	The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period. Please check the App- V Client Operational logs in the Event Log Viewer for more details.

3.2.3 App-V Client Virtual Application Log Test

This test reports the statistical information about the virtual application events generated by the target system.

Note:

This test will report metrics only when the App-V Client is installed on the Microsoft RDS Server.

Target of the test : An App-V Client on the target Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the App-V Client that is to be monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. By default, this is *8080*.

Parameters	Description
Logtype	Refers to the type of event logs to be monitored. The default value is **Microsoft-AppV-Client/Virtual Applications**.
Policy Based Filter	Using this page, administrators can configure the event sources, event IDs, and event descriptions to be monitored by this test. In order to enable administrators to easily and accurately provide this specification, this page provides the following options: • Manually specify the event sources, IDs, and descriptions in the Filter text area, or, • Select a specification from the predefined filter policies listed in the Filter box For explicit, manual specification of the filter conditions, select the **No** option against the Policy Based Filter field. This is the default selection. To choose from the list of pre-configured filter policies, or to create a new filter policy and then associate the same with the test, select the **Yes** option against the Policy Based Filter field.
Filter	If the Policy Based Filter flag is set to **No**, then a Filter text area will appear, wherein you will have to specify the event sources, event IDs, and event descriptions to be monitored. This specification should be of the following format: *{Displayname}:{event_ sources_ to_ be_ included}:{event_ sources_ to_ be_ excluded}: {event_ IDs_ to_ be_ included}: {event_ IDs_ to_ be_ excluded}: {event_ descriptions_ to_ be_ included}:{event_ descriptions_ to_ be_ excluded}*. For example, assume that the Filter text area takes the value, *OS_ events:all:Browse,Print:all:none:all:none*. Here: • *OS_events* is the display name that will appear as a descriptor of the test in the monitor UI; • *all* indicates that all the event sources need to be considered while monitoring. To monitor specific event sources, provide the source names as a comma-separated list. To ensure that none of the event sources are monitored, specify *none*. • Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, *Browse* and *Print* have been excluded from monitoring. Alternatively, you can use *all* to indicate that all the event sources have to be excluded from monitoring, or *none* to denote that none of the event sources need be excluded.

Parameters	Description
	• In the same manner, you can provide a comma-separated list of event IDs that require monitoring. The *all* in our example represents that all the event IDs need to be considered while monitoring. • Similarly, the *none* (following *all* in our example) is indicative of the fact that none of the event IDs need to be excluded from monitoring. On the other hand, if you want to instruct the eG Enterprise system to ignore a few event IDs during monitoring, then provide the IDs as a comma-separated list. Likewise, specifying *all* makes sure that all the event IDs are excluded from monitoring. • The *all* which follows implies that all events, regardless of description, need to be included for monitoring. To exclude all events, use *none*. On the other hand, if you provide a comma-separated list of event descriptions, then the events with the specified descriptions will alone be monitored. Event descriptions can be of any of the following forms - *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. • In the same way, you can also provide a comma-separated list of event descriptions to be excluded from monitoring. Here again, the specification can be of any of the following forms: *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. In our example however, none is specified, indicating that no event descriptions are to be excluded from monitoring. If you use *all* instead, it would mean that all event descriptions are to be excluded from monitoring. By default, the Filter parameter contains the value: all. Multiple filters are to be separated by semi-colons (;). **Note:** The event sources and event IDs specified here should be exactly the same as that which appears in the Event Viewer window. On the other hand, if the Policy Based Filter flag is set to **Yes**, then a Filter list box will appear, displaying the filter policies that pre-exist in the eG Enterprise

Parameters	Description
	system. A filter policy typically comprises of a specific set of event sources, event IDs, and event descriptions to be monitored. This specification is built into the policy in the following format: *{Policyname}: {event_ sources_ to_ be_ included}: {event_ sources_ to_ be_ excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}* To monitor a specific combination of event sources, event IDs, and event descriptions, you can choose the corresponding filter policy from the Filter list box. Multiple filter policies can be so selected. Alternatively, you can modify any of the existing policies to suit your needs, or create a new filter policy. To facilitate this, a ⊕ icon appears near the Filter list box, once the **Yes** option is chosen against Policy Based Filter. Clicking on the ⊕ icon leads you to a page where you can modify the existing policies or create a new one. The changed policy or the new policy can then be associated with the test by selecting the policy name from the Filter list box in this page.
UseWMI	The eG agent can either use **WMI** to extract event log statistics or directly parse the event logs using **event log API**s. If the UseWMI flag is **Yes**, then **WMI** is used. If not, the **event log API**s are used. This option is provided because on some Windows NT/2000 systems (especially ones with service pack 3 or lower), the use of WMI access to event logs can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the UseWMI parameter value to **No**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the UseWMI flag should always be set to 'Yes'.**
Stateless Alerts	Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the **EventLog** test. When this test captures an error event for the very first time, the eG manager will send out a **Critical** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **Critical**, but will not send out the

Parameters	Description
	details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting** capability. To enable this capability for this test, set the Stateless Alerts flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.
DDforinformation	eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the DDforinformation and DDforwarning flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the DDforinformation flag to **No**.
DDforwarning	To ensure that the test does not generate and store detailed measures for warning events, set the DDforwarning flag to **No**.
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 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
Information messages	Indicates the number of App- V Client virtual application informational events that were generated when the test was last executed.	Number	A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications. Please check the App- V Client Virtual Application logs in the Event Log Viewer for more details.
Warnings	Indicates the number of App- V Client virtual application warnings that were generated when the test was last executed.	Number	A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more applications. Please check the App- V Client Virtual Application logs in the Event Log Viewer for more details.
Error messages	Indicates the number of App- V Client virtual application error events that were generated during the last measurement period.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more applications. Please check the App- V Client Virtual Application logs in the Event Log Viewer for more details. Please check the App- V Client Virtual Application logs in the Event Log Viewer for more details.

Measurement	Description	Measurement Unit	Interpretation
Critical messages	Indicates the number of App- V Client virtual applications critical error events that were generated when the test was last executed.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of fatal/irrepairable problems in one or more applications. Please check the App- V Client Virtual Application logs in the Event Log Viewer for more details.
Verbose messages	Indicates the number of App- V Client virtual application verbose events that were generated when the test was last executed.	Number	The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period. Please check the App- V Client Virtual Application logs in the Event Log Viewer for more details.

3.3 The Remote Desktop Services Layer

The tests associated with this layer (see Figure 3.4) enable administrators to measure the health of the client to server connectivity, using metrics such as the following:

- The availability of the Microsoft RDS server and its responsiveness to client requests
- Login time to the server
- The status of file serving as seen by a Microsoft RDS client

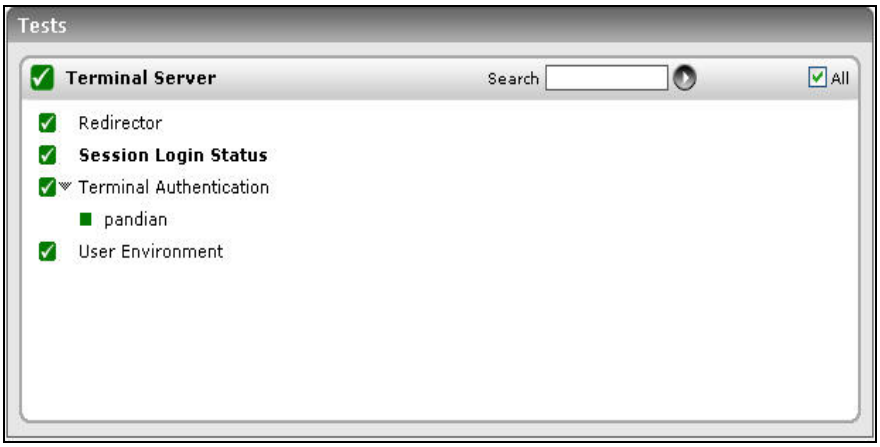

Figure 3.4: Tests associated with the Remote Desktop Services layer

3.3.1 Session Login Status Test

Administrators typically use the **Change logon command line** tool to enable / disable logons from client sessions to the Citrix / Microsoft RDS server. Disabling client logons will deny all users access to the server. Whenever users complaint of login failures, administrators might first want to check the status of the client logons to determine whether it has been disabled or not. This test periodically reports the status of logons from client sessions to the Citrix / Microsoft RDS server.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results the Microsoft RDS server 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 the test is to be configured.
Port	Enter the port to which the host listens

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Session login status	Indicates whether the client sessions to the	Percent	If the value for this measure is 100, it indicates that client logons are

Measurement	Description	Measurement Unit	Interpretation
	server are currently enabled or not.		enabled. In this case, both new client sessions and reconnections to existing sessions will be allowed. If the value of this measure is 0, it indicates that client logons are disabled. In this case, both new client sessions and reconnections to existing sessions will be disallowed. If this measure reports a value (be it 0 or 100), then the other two measures of this test will not report any value.
Are new user logons disabled ?	Indicates whether/not new user logons are disabled.		If only new user logons are disabled, then this measure will report the value Yes. This implies that, new client sessions will be disallowed, but reconnections to existing sessions will be allowed. If the value of this measure is **Yes**, then the *Session login status* measure and the Are new user logons disabled until server is restarted? measure will not report any values. If both new user logons and reconnections to existing user sessions are allowed, then this measure will not report any value. Instead, the Session login status measure will report the value 100. Likewise, if both new user logons and reconnections to existing user sessions are disallowed, then once

Measurement	Description	Measurement Unit	Interpretation
			again, this measure will not report any value. Instead, the *Session login status* measure will report the value 0. Similarly, if new user logons are not disabled permanently, but only until server restart, then, this measure will not report any value. Instead, the Are new user logons disabled until server is restarted? measure will report the value **Yes** . **Note:** In the graph of this measure, the value **Yes** for this measure will be represented using the numeric value 0.
Are new user logons disabled until server is restarted?	Indicates whether/not new user logons have been disabled until the server is restarted.		If new user logons alone are disabled until the server is restarted, then the value of this measure will be Yes. This implies that, new client sessions will be disallowed until such time the server is restarted, but reconnections to existing sessions will be allowed. If the value of this measure is **Yes**, then the Session login status measure and the Are new user logons disabled? measure will not report any values. If both new user logons and reconnections to existing user sessions are allowed, then this measure will not report any value.

Measurement	Description	Measurement Unit	Interpretation
			Instead, the *Session login status* measure will report the value 100. Likewise, if both new user logons and reconnections to existing user sessions are disallowed, then once again, this measure will not report any value. Instead, the *Session login status* measure will report the value 0. Similarly, if new user logons are disabled permanently – i.e., not just until server restart - then, this measure will not report any value. Instead, the *Are new user logons disabled?* measure will report the value *Yes* . **Note:** In the graph of this measure, the value **Yes** for this measure will be represented using the numeric value 0.

3.3.2 Terminal Connection Test

This test tracks various statistics pertaining to Microsoft RDS server connections to and from a host, from an external perspective.

Target of the test : A Microsoft RDS server

Agent deploying the test : An external agent

Outputs of the test : One set of results for every port being monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.

Parameters	Description
Host	The host for which the test is to be configured.
Port	Enter the port to which the specified TargetHost listens
TargetPorts	Specify a comma-separated list of port numbers that are to be tested (eg., 80,7077,1521). By default, the default terminal sever port, 3389, will be displayed here.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Connection availability	Whether the Microsoft RDS server 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.
Connection 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.

3.3.3 Terminal Authentication Test

This test emulates the user login process at the system level on a Microsoft RDS server and reports whether the login succeeded and how long it took.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every user account being checked

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server .
Username	This test emulates the user login process at the system level on a Microsoft RDS server. Therefore, specify the login name of a user with both **interactive logon** and **logon locally** privileges.
Password	Enter the password that corresponds to the specified Username.
Confirm Password	Confirm the password by retyping it here.
Domain	Specify the name of the domain to which the test will try to login. If the test is to login to a local host, specify 'none' here. **Note:** If users are spread across multiple domains, then, you can configure this test with multiple domain specifications; in this case, for every domain, a User-Password pair might also have to be configured. Sometimes, you might want the test to login as specific users from the same domain, to check how long each user login takes. Both these scenarios require the configuration of multiple domains and/ormultiple User names and Passwords. In order to enable users to specify these details with ease, eG Enterprise provides a special page; to access this page, click on the Click here hyperlink at the top of the parameters in the test configuration page. To know how to use this page, refer to the **Configuring Multiple Users for the Citrix Authentication Test** section in the *Monitoring Citrix XenApp Server* document.
Report by Domain Name	By default, this flag is set to **Yes**. This implies that by default, the detailed diagnosis of this test will display the *domainname\username* of each user who logged into the Microsoft RDS server. This way, administrators will be able to quickly determine which user logged in from which domain. If you want the detailed diagnosis to display the *username* alone, then set this flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Authentication status	Indicates whether the login was successful or not	Percent	A value of 100 % indicates that the login has succeeded. The value 0 is indicative of a failed login.
Authentication time	Indicates the time it took to login	Secs	If this value is very high then it could be owing to a configuration issue (i.e. the domain might not be configured properly) or a slow-down/unavailability of the primary domain server.

3.3.4 Redirector Test

File serving very often is a much underestimated part of Citrix and Microsoft RDS server environments. Improperly configured file serving components can wreak havoc on a server farm's performance.

File serving in Citrix and Microsoft RDS server environments is used at different times. For instance, every time a user logs on or off, profile data may be copied back and forth between the file server and terminal or Citrix server. Another example involves multiple applications accessing configurations stored in files from a remote file server. Folder redirection, if used, is another form of file retrievals from file servers.

File serving problems can have a detrimental impact on the performance of Citrix/Microsoft RDS server environments. Often, these problems may manifest in many ways. For example, users may see very slow access to their home directory, or folders. Even with a small profile, logging on and off could take a long time. Random application crashes can also happen, especially for applications that rely on file servers to store their configuration files remotely. Such file serving problems are often the most difficult to diagnose.

The Redirector component of the Microsoft Windows operating system handles file serving at the client end, and the Redirector test monitors this component's activity, and tracks the status of file serving as seen by a file server's client (i.e., the Citrix or Microsoft RDS server).

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the Microsoft RDS server 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 the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server .

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Data received	This metric shows the rate of data that were received by the local server from the network. This includes all the application data as well as network protocol information.	MB/Sec	
Data sent	This metric represents the rate at which data is leaving the Redirector to the network. This includes all the application data as well as network protocol information.	MB/sec	
Current commands	This metric indicates the number of requests to the Redirector that are currently queued for	Number	The *Current Commands* measure indicates the number of pending commands from the local computer to all destination servers. This

Measurement	Description	Measurement Unit	Interpretation
	service.		means that if one of the destination servers does not respond in a timely manner, the number of current commands on the local computer may increase. If the local computer is serving many sessions, a high number of current commands does not necessarily indicate a problem or a bottleneck. However, if the *Current Commands* measure shows a high number and the local computer is idle, this may indicate a network-related problem or a redirector bottleneck on the local computer. For example, there may be a network-related problem or a local bottleneck if the computer is idle overnight but the counter shows a high number during that period.
Network errors	This metric denotes the rate at which serious unexpected errors are occurring during file system access from a remote server.	Errors/sec	Such errors generally indicate that the Redirector and one or more Servers are having serious communication difficulties. For example an SMB (Server Manager Block) protocol error is a Network Error. An entry is written to the System Event Log and provides details.
Reads denied	This metric denotes the rate at which the server is unable to accommodate requests for raw read operations.	Reads/sec	When a read is much larger than the server's negotiated buffer size, the Redirector requests a Raw Read which, if granted, would permit the transfer of the data without lots of protocol overhead

Measurement	Description	Measurement Unit	Interpretation
			on each packet. To accomplish this, the server must lock out other requests, so the request is denied if the server is really busy.
Hung server sessions	This metric shows the number of active sessions that are timed out and unable to proceed due to a lack of response from the remote file server.	Number	
Writes denied	This metric denotes the rate at which the server is unable to accommodate requests for raw write operations	Writes/sec	When a write is much larger than the server's negotiated buffer size, the Redirector requests a Raw Write which, if granted, would permit the transfer of the data without lots of protocol overhead on each packet. To accomplish this, the server must lock out other requests, so the request is denied if the server is really busy.

3.3.5 User Profile Test

User profiles are the heart of the Microsoft RDS server environment. User profiles contain the configuration settings, which bring the user desktop alive. One of the major problems in a server-based computing environment like the Microsoft RDS server is that the user's login process takes more time to open the user's desktop. This happens if the user profile size is huge. The **User Profile** test monitors the size of the Microsoft RDS server user profiles and raises an alarm if the profile size exceeds the profile quota size.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every user profile on the Microsoft RDS server monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server .
Profilesizelimit	Specify the profile quota size (in MB). The default value is 50 MB.
Exclude	Provide a comma-separated list of users who need to be excluded from the analysis. By default, this parameter is set to *All_Users*, indicating that, by default, the test will not monitor the *All_Users* profile
CurrentUsersOnly	If this is set to **true**, then the profile sizes of only those users who are currently logged into the server will be monitored. If this is set to **false**, eG Enterprise will perform profile monitoring for all the users to the server.
FileSizeLimit	Takes the file quota size (in KB). The default size is 10000 KB.
Report By Domain	By default, this flag is set to **Yes**. This implies that by default, this test will report metrics for every *domainname\username* to the server. This way, administrators will be able to quickly determine which user belongs to which domain. If you want the test to report metrics for every username alone, then set this flag to *No*.
User Profile DIR	By default, this parameter is set to *none*. This implies that for XenApp/Microsoft RDS servers operating on Windows 2008 and Windows 2012 platforms, the test will, by default, check the C:\Users directory for the user profile files. In some environments, the user profile-related files and folders may exist in a different directory. In such environments, you will have to specify the exact directory in which the user profiles exist, against the User Profile DIR parameter.
Exclude Folders	By default, when this test computes the size of a profile, it automatically excludes the following folders and their sub-folders from the computation: AppData\Local,AppData\LocalLow,Recycle.Bin,SkyDrive,WorkFolders. If need be, you can choose to include one/more of these default folders when computing the profile size; for this, all you need to do is remove those specific folders from the default Exclude Folders specification. For example, to include the **SkyDrive** and **WorkFolders**

Parameters	Description
	folders, simply remove them from the default specification above. Also, if required, you can exclude more folders from the profile size computation, by appending the corresponding folder names / folder name patterns to this default list. For instance, your specification can be: **AppData\Local,AppData\LocalLow,Recycle.Bin,SkyDrive,WorkFolders,*Backup*,Favo*,*Desktop**. In the case of this sample specification, in addition to the default list of excluded folders, all folders with names that embed the string **Backup**, with names that begin with the string **Favo**, and with names that end with the string **Desktop**, will be excluded from size computation. Moreover, all sub-folders within these folders will also be ignored during size computation.
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
Is user profile exceeding quota?	Indicates whether the profile size exceeds the profile quota size by comparing the current profile size with the configured ProfileSizeLimit parameter.	Boolean	If this measure shows 0, it indicates that the current profile size has not exceeded the quota size. The value 1 indicates that the current profile size has exceeded the quota size.
Current profile size	Indicates the current	MB	

Measurement	Description	Measurement Unit	Interpretation
	profile size.		
Number of files in user's profile	Indicates the number of files available in the user profile.	Number	
Large files in user's profile	The number of files in the user profile, which exceed the allowable FileSizeLimit parameter.	Number	The detailed diagnosis of this measure, if enabled, lists all the files that have exceeded the configured FileSizeLimit.

3.3.6 Windows User Logon Test

The process of a user logging into a Citrix or Microsoft RDS server is fairly complex. First, the profile corresponding to a user has to be located, and the appropriate profile files copied over from a profile server (in the case of a roaming profile). Second, additional processing is often necessary after copying the profile locally. Processing for instance may involve creating new printers for the user logging in. Proper monitoring of profile loading and processing times is key because the login process is handled exclusively by Microsoft Windows. Hence, if a specific user profile takes a lot of time to load (e.g., because the profile is very big), or if specific processing for a user is taking time, this could delay logins for subsequent users who are trying to access the server at the same time. The typical process for monitoring the Windows login process is to use the user environment debugging mechanism. To enable this, the following steps are required. To set the logging level associated with the userenv.log file, perform the following steps:

- Start a registry editor (e.g., regedit.exe).
- Navigate to the HKEY_ LOCAL_ MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon registry subkey.
- From the **Edit** menu, select New, **DWORD** Value.
- Enter the name **UserEnvDebugLevel**, then press **Enter**.
- Double-click the new value, set it to 65538 (decimal) - which corresponds to the debugger output.

Once these changes are enabled, details about the Windows login process are logged into the file %systemroot%\debug\usermode\userenv.log. If the Userenv.log file is larger than 300 KB, the file is renamed Userenv.bak, and a new Userenv.log file is created. This action occurs when a user logs on locally or by using Terminal Services, and the Winlogon process starts. However, because the

size check only occurs when a user logs on, the Userenv.log file may grow beyond the 300 KB limit. The 300 KB limit cannot be modified.

This test periodically checks the userenv log file to monitor the user login and profile loading process. This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick the *Microsoft RDS* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the **<** button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every Microsoft RDS server monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server .
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
Profile load starts	Indicates the number of profile loads in the last measurement period.	Number	This metric gives an idea of the rate at which users are logging in to the server. The detailed diagnosis will not be available for the 'Total' descriptor.
Profile load successes	Indicates the number of successful profile loads in the last measurement period.	Number	
Profile loading failures	Indicates the number of profile load failures in the last measurement period.	Number	An unusual increase in number of profile loading failures is a cause for concern. The userenv.log file will have details of what profile loads failed and why.
Profile load failures percent	Indicates the percentage of profile loads that failed in the last measurement period.	Percent	
Avg user profile load time	Indicates the average time it took to load a profile successfully in the last measurement period.	Seconds	The detailed diagnosis of this measure, if enabled, lists the profile load times for different Citrix users.
Max profile load time	Indicates the maximum time it took to load a profile during the last measurement period.	Seconds	
System policy starts	Indicates the number of system policy applications started in the last measurement period.	Number	

Measurement	Description	Measurement Unit	Interpretation
System policy completes	Indicates the number of system policy completions in the last measurement period.	Number	Compare the total number of starts to completions. if there is a significant discrepancy, this denotes a bottleneck in system policy application. Check the userenv.log file for more details.
Avg system policy processing time	Indicates the average time taken for applying system policies in the last measurement period.	Seconds	If the system policy times are long, check the detailed diagnosis to view if the policy handling is taking time for all users. Analyze the userenv.log to determine the reason for any slowdown.
Max system policy time	Indicates the maximum time for applying system policies in the last measurement period.	Seconds	
Group policy starts	Indicates the number of group policy applications started in the last measurement period.	Number	
Group policy completes	Indicates the number of group policy applications completed in the last measurement period.	Number	
Avg group policy processing time	Indicates the average time taken for applying group policies.	Seconds	
Max group policy time	Indicates the average time taken for applying group policies.	Seconds	
Profile unload starts	Indicates the number of profile unloads started during the last	Number	The detailed diagnosis will not be available for the 'Total' descriptor.

Measurement	Description	Measurement Unit	Interpretation
	measurement period.		
Profile unload successes	Indicates the number of successful profile unloads during the last measurement period.	Number	
Profile unload failures	Indicates the number of unsuccessful profile unloads during the last measurement period.	Number	
Profile unload failures percent	Indicates the profile unload failures as a percentage of the total profile unloads.	Percent	
Avg user profile unload time	Indicates the average time for unloading a profile during the last measurement period.	Seconds	
Max profile unload time	Indicates the maximum time for unloading a profile during the last measurement period.	Seconds	
Logon duration	Indicates the average time taken by this user for logging in during the last measurement period.	Seconds	If this value is abnormally high for any user, then, you can compare the User account discovery, LDAP bind time to active directory, Client side extension processed time, DC discovery time, Total group policy object file accessed time, Avg system policy processing time and User profile load time measures to know exactly where that user's login process experienced a bottleneck - is it when loading the

Measurement	Description	Measurement Unit	Interpretation
			profile? is it when processing system policies? is it when processing group policies? is it when interacting with AD for authenticating the user login? This measure will not be available for Citrix XenApp Servers operating on Windows 2003.
User account discovery	Indicates the amount of time taken by the system call to get account information for this user during the last measurement period.	Seconds	Compare the value of this measure across users to know which user's logon process spent maximum time in retrieving account information. The detailed diagnosis will not be available for the 'Total' descriptor. This measure will not be available for Citrix XenApp Servers operating on Windows 2003.
LDAP bind time to active directory	Indicates the amount of time taken by the LDAP call for this user to connect and bind to Active Directory during the last measurement period.	Seconds	Compare the value of this measure across users to know which user's logon process spent maximum time in connecting to Active Directory. Besides impacting authentication time, high LDAP bind time may also affect group policy processing. This measure will not be available for Citrix XenApp Servers operating on Windows 2003.
Client side extensions applied	Indicates the number of client side extensions used for processing group policies for this user during the last measurement period.	Number	

Measurement	Description	Measurement Unit	Interpretation
Client side extension processed time	Indicates the amount of time that client side extensions took for processing group policies for this user during the last measurement period.	Seconds	Compare the value of this measure across users to know which user's logon process spent maximum time in group policy processing. If this measure reports an unusually high value for any user, then, you may want to check the value of the LDAP bind time to active directory measure for that user to figure out if a delay in connecting to AD is affecting group policy processing. This is because, group policies are built on top of AD, and hence rely on the directory service's infrastructure for their operation. As a consequence, DNS and AD issues may affect Group Policies severely. One could say that if an AD issue does not interfere with authentication, at the very least it will hamper group policy processing. You can also use the detailed diagnosis of this measure to know which client side extension was used to process which group policy for a particular user. The detailed diagnosis will not be available for the 'Total' descriptor. This measure will not be available for Citrix XenApp Servers operating on Windows 2003.
DC discovery time	Indicates the time taken to discover the domain	Seconds	Compare the value of this measure across users to know which user's

Measurement	Description	Measurement Unit	Interpretation
	controller to be used for processing group policies for this user during the last measurement period.		logon process spent maximum time in domain controller discovery. The detailed diagnosis will not be available for the 'Total' descriptor. This measure will not be available for Citrix XenApp Servers operating on Windows 2003.
Total group policy object file accessed time	Indicates the amount of time the logon process took to access group policy object files for this user during the last measurement period.	Seconds	Compare the value of this measure across users to know which user's logon process spent maximum time in accessing the group policy object file. This measure will not be available for Citrix XenApp Servers operating on Windows 2003.
User profile load time	Indicates the amount of time it took to load this user's profile successfully in the last measurement period.	Seconds	Compare the value of this measure across users to know which user's profile took the longest time to load. One of the common reasons for long profile load times is large profile size. In such circumstances, you can use the User Profile test to determine the current size of this user's profile. If the profile size is found to be large, you can conclude that it is indeed the size of the profile which is affecting the profile load time. Another reason would be the absence of a profile. If the user does not already have a profile a new one is created. This slows down the initial logon quite a bit compared to subsequent logons.

Measurement	Description	Measurement Unit	Interpretation
			The main reason is that Active Setup runs the IE/Mail/Theme initialization routines. 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. This measure will not be available for Citrix XenApp Servers operating on Windows 2003.

3.3.7 RDS CALs Test

This test reports the usage statistics pertaining to a Microsoft RDS server's client access licenses. To ensure that the test functions smoothly, the Terminal Services Licensing Reporter tool (*lsreport.exe*) needs to be available on the eG agent host. *lsreport.exe* is a command-line utility that you can use to display information about the licenses that are issued by Microsoft RDS License servers. *lsreport.exe* connects to Microsoft RDS License servers and logs information about the license key packs that are installed on the servers. In order to make sure that this utility is available to the eG Enterprise suite, do the following:

- Download the *lsreport.exe* from the Microsoft Windows 2000 Server Resource Kit.
- Copy *lsreport.exe* to the {EG_INSTALL_DIR}\bin directory.

This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick the *Microsoft RDS* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the **<** button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the Microsoft RDS server 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 the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server .
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
Active licenses	Represents number of active client access licenses that were currently consumed from the Microsoft RDS server license server.	Number	The detailed diagnosis of this provides the complete details of the active access licenses, which include critical session information such as the user who initiated the session, the start and end date/time of the session, the type of license issued to the user, the license ID, the issue type, the

Measurement	Description	Measurement Unit	Interpretation
			target server, the client from which the session was instantiated, etc.
Temporary licenses	Indicates the number of temporary client access licenses that were currently consumed from the Microsoft RDS server license.	Number	The detailed diagnosis of this provides the complete details of the temporary access licenses, which include critical session information such as the user who initiated the session, the start and end date/time of the session, the type of license issued to the user, the license ID, the issue type, the target server, the client from which the session was instantiated, etc.

3.3.8 GDI Objects Test

An object is a data structure that represents a system resource, such as a file, thread, or graphic image. An application cannot directly access object data or the system resource that an object represents. Instead, an application must obtain an object handle, which it can use to examine or modify the system resource. There are three categories of objects: user, GDI, and kernel. GDI objects support graphics. Here is a list of the GDI objects used in Windows:

- Bitmap
- Brush
- Device Context (DC)
- Enhanced Metafile
- Enhanced-metafile DC
- Font
- Memory DC
- Metafile
- Metafile DC
- Palette

- Pen/extended pen
- Region

GDI objects support only one handle per object, and only the process that created the object can use the object handle.

If an application creates a lot of these objects, without properly destroying references to the object (by closing the associated handle), then there will be multiple GDI objects occupying memory on the system for each object created. If this GDI leak is really bad, this can eventually bring a server to its knees, and cause all types of problems (slow logons, registry issues, system hangs, and so on).

If such fatalities are to be avoided, administrators should closely monitor the number of GDI object handles created by every user to the Microsoft RDS server and proactively detect potential GDI leaks. This is where the **GDI Objects** test helps. This test periodically checks the GDI object handles created by each user to the Microsoft RDS server, reports the total number of handles created per user, and promptly notifies administrators if any user is creating more GDI handles than permitted. This way, the test brings probable GDI leaks to the attention of administrators. In addition, administrators can use the detailed diagnosis of the test to know which process is responsible for the GDI leak (if any).

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each user to the Microsoft RDS server 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 the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server.
GDILimit	Specify the maximum number of GDI object handles that a user to the Microsoft RDS server can create. By default, this value is 10000.
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,

Parameters	Description
	click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total GDI objects	Indicates the total number of GDI handles that this user has created.	Number	The detailed diagnosis of this measure, if enabled, provides the process-wise breakup of the GDI handles created by the user. In the event of a GDI leak, this information will enable you to figure out which process initiated by the user spawned the maximum number of GDI handles, and is hence responsible for the GDI leak.
Percentage of GDI objects	Indicates what percentage of the configured gdilimit is the total number of GDI object handles created by this user's processes.	Percent	This value is calculated using the following formula: **Total GDI objects/GDILimit * 100** A value close to 100% is a cause for concern, as it indicates that the count of GDI handles for the user is fast- approaching the permitted gdilimit. This hints at a potential GDI leak. You can then use the detailed diagnosis of the Total GDI objects measure to identify which process initiated by the user is spawning the maximum GDI handles and is

Measurement	Description	Measurement Unit	Interpretation
			hence contributing to the leak, and probe further.

The detailed diagnosis of the *Total GDI* objects measure, if enabled, provides the process-wise breakup of the GDI handles created by the user. In the event of a GDI leak, this information will enable you to figure out which process initiated by the user spawned the maximum number of GDI handles, and is hence responsible for the GDI leak.

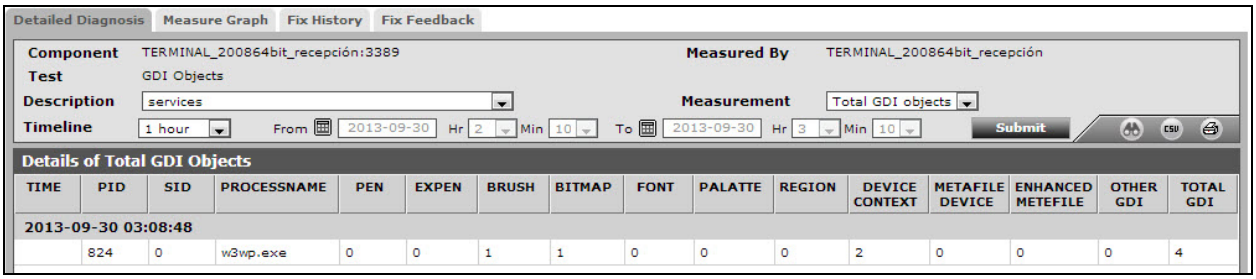

Figure 3.5: The detailed diagnosis of the Total GDI objects measure

3.3.9 ICA/RDP Listeners Test

The listener component runs on the XenApp/Microsoft RDS server and is responsible for listening for and accepting new ICA/RDP client connections, thereby allowing users to establish new sessions on the XenApp/Microsoft RDS server. If this listener component is down, users may not be able to establish a connection with the XenApp server!

This is why, if a user to the Microsoft RDS server complains of the inaccessibility of the server, administrators should first check whether the listener component is up and running or not. The **ICA/RDP Listeners** test helps administrators perform this check. This test tracks the status of the default listener ports and reports whether any of the ports is down.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of outputs for every listener port configured

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.

Parameters	Description
Host	The host for which the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server.
Session IDs	The default listener ports - *65536,65537,65538* – will be displayed here by default. You can override this default specification by adding more ports or by removing one/more existing ports.

Measurements of the test

Measurement	Description	Measurement Unit	Interpretation
Is listener down?	Indicates whether/not this listener port is down.		This measure reports the value Yes if the listener port is down and No if the port is up and running. The numeric values that correspond to these measure values are as follows: Measure Value: Yes — Numeric Value: 0 Measure Value: No — Numeric Value: 1 **Note:** By default, this measure reports the above-mentioned **Measure Value**s to indicate the status of a listener port. However, the graph of this measure will represent the same using the numeric equivalents only.

3.3.10 User Logon Details Test

The process of a user logging into a Microsoft RDS server is fairly complex. First, the domain controller is discovered and the login credentials are authenticated. Then, the corresponding user profile is identified and loaded. Next, group policies are applied and logon scripts are processed to setup the user environment. In the meantime, additional processing may take place for a user – say,

applying system profiles, creating new printers for the user, and so on. A slowdown in any of these steps can significantly delay the logon process for a user. Since logons on Windows happen sequentially, this may adversely impact the logins for other users who may be trying to access the Microsoft RDS server at the same time. Hence, if a user complains that he/she is unable to access an application/desktop published on Microsoft RDS, administrators must be able to rapidly isolate exactly where the logon process is stalling and for which user.

This test periodically monitors the user login and profile loading process and accurately identify where the process is bottlenecked. This test helps administrators to capture anomalies in the user login and profile loading process and report where the process is bottlenecked - in the authentication process? during profile loading? during GPO processing and if so, which GPO? which client side extension was processed by the GPO when the delay occurred? is the group policy processing mode impacting user logon?

By default, this test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick the *Microsoft RDS* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the **<** button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every user to the Microsoft RDS server monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	Refers to the port at which the host listens.
Report Total	By default, this flag is set to **No**. In this case therefore, the test will only report metrics for every user to the target server. If this flag is set to **Yes**, then the test will report metrics for a *Total* descriptor - the metrics reported by this descriptor will be aggregated across all users to the target server. This way, the administrators will receive a system-wide overview of the health of the profile loading/unloading process.
Profilesizelimit	Specify the profile quota size (in MB). The default value is 50 MB.

Parameters	Description
FileSizeLimit	Takes the file quota size (in KB). The default size is 10000 KB.
Exclude Folders	By default, when this test computes the size of a profile, it automatically excludes the following folders and their sub-folders from the computation: AppData\Local,AppData\LocalLow,Recycle.Bin,SkyDrive,WorkFolders. If need be, you can choose to include one/more of these default folders when computing the profile size; for this, all you need to do is remove those specific folders from the default Exclude Folders specification. For example, to include the **SkyDrive** and **WorkFolders** folders, simply remove them from the default specification above. Also, if required, you can exclude more folders from the profile size computation, by appending the corresponding folder names / folder name patterns to this default list. For instance, your specification can be: **AppData\Local,AppData\LocalLow,Recycle.Bin,SkyDrive,WorkFolders,*Backup*,Favo*,*Desktop**. In the case of this sample specification, in addition to the default list of excluded folders, all folders with names that embed the string **Backup**, with names that begin with the string **Favo**, and with names that end with the string **Desktop**, will be excluded from size computation. Moreover, all sub-folders within these folders will also be ignored during size computation.
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 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
Logon duration	Indicates the average time taken by this user for logging in during the last measurement period.	Secs	If this value is abnormally high for any user, then, you can compare the time measurements reported under Logon Phase to know where exactly the user logon was bottlenecked - was it when loading the profile? when processing group policies? when authenticating the user over the network? when initializing the user? when initializing the desktop? Use the detailed diagnosis of this measure to know which were the client side extensions processed for a user, which group policy was processed by each extension, and what is the processing time of each client side extension. In the process, you can quickly identify the client side extension that is taking an abnormally long time for processing and is probably the root-cause for the logon delay.
Network providers duration	Indicates the amount of time taken by the network provider to authenticate this user on their network.	Secs	A Network Provider is a DLL which is responsible for a certain type of connection protocol. On each logon, Winlogon notifies these Network Providers so that they can collect credentials and authenticate the users on their network. Citrix PnSson is a common network provider found on XenApp and XenDesktop VM's. If Logon duration is abnormally high, then compare the value of this measure with the other time measurements displayed under Logon Phase to determine whether/not the logon delay can be attributed to an authentication delay. Sometimes, an authentication delay can be caused by an interim delay between when

Measurement	Description	Measurement Unit	Interpretation
			the Network provider phase ends and when the next phase - i.e., the Citrix profile management phase - begins. To verify whether/not such an interim delay has occurred and to assess its impact on the network providers authentication process, use the detailed diagnosis of this measure. The detailed diagnosis reveals the time that elapsed between when the Network providers phase ended and the next phase began. Compare this interim delay with the value of this measure to understand whether/not the interim delay is the reason for the delay in the Network providers phase. **Note:** By default, this test does not report the *Network providers duration* measure. If you want the test to report this measure, then make sure that *Audit process tracking* is enabled on the managed Microsoft RDS server. The Audit process tracking setting determines whether/not to audit detailed tracking information for events such as program activation, process exit, handle duplication, and indirect object access. To enable this setting, do the following: 1. First, check whether/not the setting is already enabled. For that, login to the managed server. 2. Go to the command prompt. 3. Run the following command at the prompt: `Auditpol /get /category:*`

Measurement	Description	Measurement Unit	Interpretation

			4. If this command returns the value *No auditing*, it implies that the Audit process tracking setting is not enabled on the server.
			5. In this case, proceed to enable the setting by issuing the following commands at the prompt, one after another:
			`Auditpol /set /subcategory:"Process Creation" /success:enable`
			`Auditpol /set /subcategory:"Process Termination" /success:enable`
			`Auditpol /set /subcategory:"Process Creation" /failure:enable`
			`Auditpol /set /subcategory:"Process Termination" /failure:enable`
			6. After enabling the Audit process tracking setting, ensure that the size of the security log is increased to 100 MB. This is required because, once Audit process tracking is enabled, the security log will be flooded with numerous Process creation and Process termination events. If the security log is not sized adequately, then it will not be able to capture these events. Consequently, the test will not be able to report this measure. To avoid this, it is imperative that the

Measurement	Description	Measurement Unit	Interpretation
			security log size is increased.
Citrix profile management duration	Indicates the amount of time taken to load the citrix profile of this user successfully during the last measurement period.	Secs	During logon, Citrix UPM copies the users' registry entries and files from the user store to the local profile folder. If a local profile cache exists, the two sets are synchronized.
User profile duration	Indicates the amount of time it took to load this user's profile successfully in the last measurement period.	Secs	If Logon duration is abnormally high, then compare the value of this measure with the other time measurements displayed under Logon Phase to determine whether/not the logon delay can be attributed to a delay in user profile loading. One of the common reasons for long profile load times is large profile size. In such circumstances, you can use the User Profile test to determine the current size of this user's profile. If the profile size is found to be large, you can conclude that it is indeed the size of the profile which is affecting the profile load time. Another reason would be the absence of a profile. If the user does not already have a profile a new one is created. This slows down the initial logon quite a bit compared to subsequent logons. The main reason is that Active Setup runs the IE/Mail/Theme initialization routines. Sometimes, profile loading can be delayed by an interim delay between when the User profile phase ends and when the next phase - i.e., the Group policy processing phase - begins. To verify whether/not such an interim delay has occurred and to assess its impact on the profile loading process, use the detailed diagnosis of this measure. The detailed diagnosis reveals

Measurement	Description	Measurement Unit	Interpretation
			the time that elapsed between when the User profile phase ended and the next phase began. Compare this interim delay with the value of this measure to understand whether/not the interim delay is the reason for the delay in profile loading. **Note:** By default, this test does not report the *User profile duration* measure. If you want the test to report this measure, then make sure that *Audit process tracking* is enabled on the managed Microsoft RDS server. The Audit process tracking setting determines whether/not to audit detailed tracking information for events such as program activation, process exit, handle duplication, and indirect object access. To enable this setting, do the following: 1. First, check whether/not the setting is already enabled. For that, login to the managed server. 2. Go to the command prompt. 3. Run the following command at the prompt: `Auditpol /get /category:*` 4. If this command returns the value *No auditing*, it implies that the Audit process tracking setting is not enabled on the server. 5. In this case, proceed to enable the setting by issuing the following commands at the prompt, one after another:

Measurement	Description	Measurement Unit	Interpretation

			`Auditpol /set /subcategory:"Process Creation" /success:enable`
			`Auditpol /set /subcategory:"Process Termination" /success:enable`
			`Auditpol /set /subcategory:"Process Creation" /failure:enable`
			`Auditpol /set /subcategory:"Process Termination" /failure:enable`
			6. After enabling the Audit process tracking setting, ensure that the size of the security log is increased to 100 MB. This is required because, once Audit process tracking is enabled, the security log will be flooded with numerous Process creation and Process termination events. If the security log is not sized adequately, then it will not be able to capture these events. Consequently, the test will not be able to report this measure. To avoid this, it is imperative that the security log size is increased.

Group policy processing time	Indicates the time taken for applying group policies for this user in the last measurement period.	Secs	Group policies impact logon performance. If Logon duration is abnormally high, then compare the value of this measure with the other time measurements displayed under Logon Phase to determine whether/not the logon delay can be attributed to a delay in group policy processing.

Measurement	Description	Measurement Unit	Interpretation
			Typically, group policies can be processed in the foreground or background. While background processing of group policies have little to no impact on logon performance, foreground group policy processing can cause logon delays. Particularly, if foreground processing is done synchronously, the user's logon experience is bound to suffer. This is because,when foreground processing is synchronous, the user is not presented with the logon prompt until computer GP processing has completed after a system boot. Likewise the user will not see their desktop at logon until user GP processing completes. This can increase startup time as seen by the user. Other factors that can impact group policy processing time and consequently, the logon time are, long-running WQL queries and group policy scripts. Sometimes, group policy processing can be delayed by an interim delay between when the Group policy phase ends and when the next phase - i.e., the Group policy script execution phase - begins. To verify whether/not such an interim delay has occurred and to assess its impact on group policy processing, use the detailed diagnosis of this measure. The detailed diagnosis reveals the time that elapsed between when the Group policy phase ended and the next phase began. Compare this interim delay with the value of this measure to understand whether/not the interim delay is the reason for the delay in group policy processing. **Note:**

Measurement	Description	Measurement Unit	Interpretation

			By default, this test does not report the *Group policy processing time* measure. If you want the test to report this measure, then make sure that *Audit process tracking* is enabled on the managed Microsoft RDS server. The Audit process tracking setting determines whether/not to audit detailed tracking information for events such as program activation, process exit, handle duplication, and indirect object access. To enable this setting, do the following:
			1. First, check whether/not the setting is already enabled. For that, login to the managed server.
			2. Go to the command prompt.
			3. Run the following command at the prompt:
			`Auditpol /get /category:*`
			4. If this command returns the value *No auditing*, it implies that the Audit process tracking setting is not enabled on the server.
			5. In this case, proceed to enable the setting by issuing the following commands at the prompt, one after another:
			`Auditpol /set /subcategory:"Process Creation" /success:enable`
			`Auditpol /set /subcategory:"Process Termination" /success:enable`

Measurement	Description	Measurement Unit	Interpretation

			`Auditpol /set /subcategory:"Process Creation" /failure:enable`
			`Auditpol /set /subcategory:"Process Termination" /failure:enable`
			6. After enabling the Audit process tracking setting, ensure that the size of the security log is increased to 100 MB. This is required because, once Audit process tracking is enabled, the security log will be flooded with numerous Process creation and Process termination events. If the security log is not sized adequately, then it will not be able to capture these events. Consequently, the test will not be able to report this measure. To avoid this, it is imperative that the security log size is increased.

Group policy scripts duration	Indicates the time taken for executing group policy scripts for this user in the last measurement period.	Secs	If Logon duration is abnormally high, then compare the value of this measure with the other time measurements displayed under Logon Phase to determine whether/not the logon delay can be attributed to slow group policy script execution. If logon scripts are configured to run synchronously, they can cause logon delays. This is because, if logon scripts run synchronously, then the system will have to wait for the logon scripts to finish running before it starts the Windows Explorer Interface program and creates the desktop. This will delay the appearance of the desktop, which in turn, will delay user

Measurement	Description	Measurement Unit	Interpretation
			logon. Sometimes, group policy script execution can be delayed by an interim delay between when the Group policy script phase ends and when the next phase - i.e., the Pre-shell phase - begins. To verify whether/not such an interim delay has occurred and to assess its impact on script execution, use the detailed diagnosis of this measure. The detailed diagnosis reveals the time that elapsed between when the Group policy script phase ended and the next phase began. Compare this interim delay with the value of this measure to understand whether/not the interim delay has significantly delayed is the reason for the delay in group policy script execution. **Note:** By default, this test does not report the *Group policy script duration* measure. If you want the test to report this measure, then make sure that *Audit process tracking* is enabled on the managed Microsoft RDS server. The Audit process tracking setting determines whether/not to audit detailed tracking information for events such as program activation, process exit, handle duplication, and indirect object access. To enable this setting, do the following: 1. First, check whether/not the setting is already enabled. For that, login to the managed server. 2. Go to the command prompt. 3. Run the following command at the prompt:

Measurement	Description	Measurement Unit	Interpretation

			`Auditpol /get /category:*`
			4. If this command returns the value *No auditing*, it implies that the Audit process tracking setting is not enabled on the server.
			5. In this case, proceed to enable the setting by issuing the following commands at the prompt, one after another:
			`Auditpol /set /subcategory:"Process Creation" /success:enable`
			`Auditpol /set /subcategory:"Process Termination" /success:enable`
			`Auditpol /set /subcategory:"Process Creation" /failure:enable`
			`Auditpol /set /subcategory:"Process Termination" /failure:enable`
			6. After enabling the Audit process tracking setting, ensure that the size of the security log is increased to 100 MB. This is required because, once Audit process tracking is enabled, the security log will be flooded with numerous Process creation and Process termination events. If the security log is not sized adequately, then it will not be able to capture these events. Consequently, the test will not be able to report this measure.

Measurement	Description	Measurement Unit	Interpretation
			To avoid this, it is imperative that the security log size is increased.
Pre-Shell duration	Indicates the time taken to execute *Userinit.exe* for this user during the last measurement period.	Secs	The Winlogon service runs *Userinit.exe*, which runs logon scripts, reestablishes network connections, and then starts *Explorer.exe*, the Windows user interface. On RDSH sessions, *Userinit.exe* also executes the Appsetup entries such as *cmstart.exe* which in-turn calls *wfshell.exe*. If Logon duration is abnormally high, then compare the value of this measure with the other time measurements displayed under Logon Phase to determine whether/not the logon delay can be attributed to slowness in the completion of user initialization tasks. Sometimes, pre-shell tasks (i.e., user initialization tasks) can be delayed by an interim delay between when the Pre-shell phase ends and when the next phase - i.e., the Shell phase - begins. To verify whether/not such an interim delay has occurred and to assess its impact on pre-shell duration, use the detailed diagnosis of this measure. The detailed diagnosis reveals the time that elapsed between when the Pre-shell phase ended and the next phase began. Compare this interim delay with the value of this measure to understand whether/not the interim delay is the reason for the delay in pre-shell tasks. **Note:** By default, this test does not report the *Pre-shell duration* measure. If you want the test to report this measure, then make

Measurement	Description	Measurement Unit	Interpretation

			sure that *Audit process tracking* is enabled on the managed Microsoft RDS server. The Audit process tracking setting determines whether/not to audit detailed tracking information for events such as program activation, process exit, handle duplication, and indirect object access. To enable this setting, do the following:
			1. First, check whether/not the setting is already enabled. For that, login to the managed server.
			2. Go to the command prompt.
			3. Run the following command at the prompt:
			`Auditpol /get /category:*`
			4. If this command returns the value *No auditing*, it implies that the Audit process tracking setting is not enabled on the server.
			5. In this case, proceed to enable the setting by issuing the following commands at the prompt, one after another:
			`Auditpol /set /subcategory:"Process Creation" /success:enable`
			`Auditpol /set /subcategory:"Process Termination" /success:enable`
			`Auditpol /set /subcategory:"Process Creation" /failure:enable`

Measurement	Description	Measurement Unit	Interpretation
			`Auditpol /set /subcategory:"Process Termination" /failure:enable` 6. After enabling the Audit process tracking setting, ensure that the size of the security log is increased to 100 MB. This is required because, once Audit process tracking is enabled, the security log will be flooded with numerous Process creation and Process termination events. If the security log is not sized adequately, then it will not be able to capture these events. Consequently, the test will not be able to report this measure. To avoid this, it is imperative that the security log size is increased.
Shell duration	Indicates the time interval between the beginning of desktop initialization and the time the desktop became available to this user including the Active Setup Phase. Active Setup is a mechanism for executing commands once per user early during login. Active Setup is used by some operating system components like Internet Explorer to set up an initial configuration for new users logging on for	Secs	If Logon duration is abnormally high, then compare the value of this measure with the other time measurements displayed under Logon Phase to determine whether/not the logon delay can be attributed to a delay in desktop initialization. Sometimes, shell tasks (i.e., desktop initialization tasks) can be delayed by an interim delay between when the Shell phase ends and when the next phase begins. To verify whether/not such an interim delay has occurred and to assess its impact on the time taken by Shell tasks, use the detailed diagnosis of this measure. The detailed diagnosis reveals the time that elapsed between when the Shell phase ended and the next phase began. Compare this interim delay with the value of this

Measurement	Description	Measurement Unit	Interpretation

	the first time.		measure to understand whether/not the interim delay is the reason for the delay in Shell tasks.
			Note:
			By default, this test does not report the *Shell duration* measure. If you want the test to report this measure, then make sure that *Audit process tracking* is enabled on the managed Microsoft RDS server. The Audit process tracking setting determines whether/not to audit detailed tracking information for events such as program activation, process exit, handle duplication, and indirect object access. To enable this setting, do the following:
			1. First, check whether/not the setting is already enabled. For that, login to the managed server.
			2. Go to the command prompt.
			3. Run the following command at the prompt:
			`Auditpol /get /category:*`
			4. If this command returns the value *No auditing*, it implies that the Audit process tracking setting is not enabled on the server.
			5. In this case, proceed to enable the setting by issuing the following commands at the prompt, one after another:
			`Auditpol /set /subcategory:"Process Creation" /success:enable`

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<pre>Auditpol /set /subcategory:"Process Termination" /success:enable</pre>
<pre>Auditpol /set /subcategory:"Process Creation" /failure:enable</pre>
<pre>Auditpol /set /subcategory:"Process Termination" /failure:enable</pre>
6. After enabling the Audit process tracking setting, ensure that the size of the security log is increased to 100 MB. This is required because, once Audit process tracking is enabled, the security log will be flooded with numerous Process creation and Process termination events. If the security log is not sized adequately, then it will not be able to capture these events. Consequently, the test will not be able to report this measure. To avoid this, it is imperative that the security log size is increased.
</td></tr>
<tr><td>Group Policy processing status</td><td>Indicates the current status of the Group policy that is applied for this user.</td><td></td><td>The values reported by this measure and their corresponding numeric equivalents are described in the table below:
<table>
<tr><th>Measure Values</th><th>Numeric Values</th></tr>
<tr><td>Success</td><td>1</td></tr>
<tr><td>Warning</td><td>2</td></tr>
<tr><td>Error</td><td>3</td></tr>
</table>
Note:</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			By default, this measure reports the above-mentioned **Measure Value**s while indicating the current status of the Group policy. However, in the graph of this measure, the values will be represented using the corresponding numeric equivalents i.e., 1 to 3.
User account discovery	Indicates the amount of time taken by the system call to get account information for this user during the last measurement period.	Secs	Compare the value of this measure across users to know which user's logon process spent maximum time in retrieving account information. To know which domain controller and DNS is being used, use the detailed diagnosis of this measure.
LDAP bind time to active directory	Indicates the amount of time taken by the LDAP call for this user to connect and bind to Active Directory during the last measurement period.	Secs	Compare the value of this measure across users to know which user's logon process spent maximum time in connecting to Active Directory. Besides impacting authentication time, high LDAP bind time may also affect group policy processing.
Domain Controller discovery time	Indicates the time taken to discover the domain controller to be used for processing group policies for this user during the last measurement period.	Secs	Compare the value of this measure across users to know which user's logon process spent maximum time in domain controller discovery.
Total group policy object file accessed time	Indicates the amount of time the logon process took to access group policy object files for this user during the last measurement period.	Secs	Compare the value of this measure across users to know which user's logon process spent maximum time in accessing the group policy object file. To know which files were accessed and the time taken to access each file, use the detailed diagnosis of this measure. With the help of the detailed diagnostics, you can accurately isolate the object file that

Measurement	Description	Measurement Unit	Interpretation
			took the longest to access, and thus delayed the logon process.
Total client- side extensions applied	Indicates the total number of client side extensions used for processing group policies for this user during the last measurement period.	Number	
Client- side extensions with success state	Indicates the number of client side extensions that were successfully used for processing group policies for this user during the last measurement period.	Number	Use the detailed diagnosis of this measure to know which were the successful client side extensions for a user, and which group policy was processed by each extension.
Client- side extensions with warning state	Indicates the number of warnings received when client side extensions were used for processing group policies for this user during the last measurement period.	Number	Use the detailed diagnosis of this measure to know which were the client side extensions that resulted in the generation of warning events at the time of processing. You will also know which group policies were processed by each extension.
Client- side extensions with error state	Indicates the number of errors registered when client side extensions were used for processing group policies for this user during the last measurement period.	Number	Ideally, the value of this measure should be zero. A sudden/gradual increase in the value of this measure is a cause of concern. If a non-zero value is reported for this measure, then use the detailed diagnosis of this measure to know which client side extensions resulted in processing errors. You will also know which group policies were processed by each such extension. Moreover, the error code will also be displayed as part of detailed diagnostics, so that you can figure out what type of error

Measurement	Description	Measurement Unit	Interpretation
			occurred when processing the client side extensions.
Total client- side extension processed time	Indicates the total amount of time that client side extensions took for processing group policies for this user during the last measurement period.	Secs	Compare the value of this measure across users to know which user's logon process spent maximum time in client side extension processing. If this measure reports an unusually high value for any user, then, you may want to check the value of the LDAP bind time to active directory measure for that user to figure out if a delay in connecting to AD is affecting group policy processing. This is because, group policies are built on top of AD, and hence rely on the directory service's infrastructure for their operation. As a consequence, DNS and AD issues may affect Group Policies severely. One could say that if an AD issue does not interfere with authentication, at the very least it will hamper group policy processing. You can also use the detailed diagnosis of this measure to know which client side extension was used to process which group policy for a particular user. Detailed diagnostics also reveal the processing time for each client side extension. This way, you can quickly identify the client side extension that took too long to be processed and thus delayed the user logon.
Estimated network bandwidth between VM and Domain Controller	Indicates the estimated network bandwidth between the VM and domain controller for this user during the last measurement period.	Kbps	
Is bandwidth	Indicates whether/not		Several components of Group Policy rely

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>between VM and Domain Controller slow?</td><td>the network connection between the VM and domain controller is currently slow for this user.</td><td></td><td>on a fast network connection. If a fast connection is unavailable between a VM and the DOC, group policy processing can be delayed. This is why, if the Group policy processing time measure reports an abnormally high value, you may want to check the value of the Is bandwidth between VM and domain controller slow? measure to determine whether the network connection between the VM and domain controller is slow.

If the network connection between the VM and domain controller is slow for a user, then this measure will report the value Yes. If it is fast, then this measure will report the value No (connection is fast).

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

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No (connection is fast)</td><td>2</td></tr></table>
Note:

• By default, this test reports the Measure Values listed in the table above to indicate the quality of the network link between the VM and the domain controller. In the graph of this measure however, the same is indicated using the numeric equivalents only.

• To determine whether the network link is</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			slow or fast, the Group Policy service compares the result of the estimated bandwidth to the slow link threshold (configured by Group Policy). A value below the threshold results in the Group Policy service flagging the network connection as a slow link. This measure reports the status of this flag only. To know the slow link threshold that the Group Policy has configured for this link, use the detailed diagnosis of this measure.
Is the user's profile size large?	Indicates whether the profile size of this user exceeds the profile quota size by comparing the current profile size with the configured ProfileSizeLimit parameter.	Boolean	If this measure shows 0, it indicates that the current profile size has not exceeded the quota size. The value 1 indicates that the current profile size has exceeded the quota size.
Current profile size	Indicates the current profile size of this user.	MB	
Number of files in user's profile	Indicates the number of files available in this user profile.	Number	
Large files in user's profile	The number of files in this user profile, which exceed the allowable FileSizeLimit parameter.	Number	The detailed diagnosis of this measure, if enabled, lists all the files that have exceeded the configured FileSizeLimit.
Group policy applied on	Indicates whether the group policy for this user is applied during foreground processing or background processing.		Foreground and background processing are key concepts in Group Policy. Foreground processing only occurs when the machine starts up or when the user logs on. Some policy areas (also called Client Side

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>Extensions (CSEs)) can only run during foreground processing. Examples of these include Folder Redirection, Software Installation and Group Policy Preferences Drive Mapping. In contrast, background processing is that thing that occurs every 90 or so minutes on Windows workstations, where GP refreshes itself periodically. Background processing happens in the background, while the user is working and they generally never notice it. While background processing does not impact performance, foreground processing can extend start and login times.

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

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Background</td><td>1</td></tr><tr><td>Foreground</td><td>2</td></tr></table>
Note:

By default, this test reports the Measure Values listed in the table above to indicate when the group policy of a user was applied. In the graph of this measure however, the same is indicated using the numeric equivalents only.</td></tr>
<tr><td>Group policy processing mode</td><td>Indicates whether the group policies of this user are processed in the synchronous or asynchronous mode.</td><td></td><td>Foreground processing can operate under two different modes - synchronously or asynchronously. Asynchronous GP processing does not prevent the user from using their desktop while GP processing completes. For example, when the</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			computer is starting up, GP asynchronous processing starts to occur for the computer, and in the meantime, the user is presented the Windows logon prompt. Likewise, for asynchronous user processing, the user logs on and is presented with their desktop while GP finishes processing. The user is not delayed getting either their logon prompt or their desktop during asynchronous GP processing. When foreground processing is synchronous, the user is not presented with the logon prompt until computer GP processing has completed after a system boot. Likewise the user will not see their desktop at logon until user GP processing completes. This can have the effect of making the user feel like the system is running slow. In short, synchronous processing can impact startup time, where asynchronous does not. Foreground processing will run synchronously for two reasons: • The administrator forces synchronous processing through a policy setting. This can be done by enabling the Computer ConfigurationPoliciesAdministrative TemplatesSystemLogonAlways wait for the network at computer startup and logon policy setting. Enabling this setting will make all foreground processing synchronous. This is commonly used for troubleshooting problems with Group Policy processing, but does not always get turned back off again.

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
A particular CSE requires synchronous foreground processing. There are four CSEs provided by Microsoft that currently require synchronous foreground processing: Software Installation, Folder Redirection, Microsoft Disk Quota and GP Preferences Drive Mapping. If any of these are enabled within one or more GPOs, they will trigger the next foreground processing cycle to run synchronously when they are changed.
It is therefore best to avoid synchronous CSEs and to not force synchronous policy. If usage of synchronous CSEs is necessary, minimize changes to these policy settings.
The values that this measure can report and their corresponding numeric values are listed in the table below:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Synchronous</td><td>1</td></tr>
<tr><td>Asynchronous</td><td>2</td></tr>
</table>
Note:
By default, this test reports the Measure Values listed in the table above to indicate when the group policy of a user was applied. In the graph of this measure however, the same is indicated using the numeric equivalents only.
</td></tr>
</table>

3.4 The Terminal Applications Layer

The health of a Microsoft RDS server depends upon the health of the applications it hosts. The Terminal Applications test associated with this layer monitors application health.

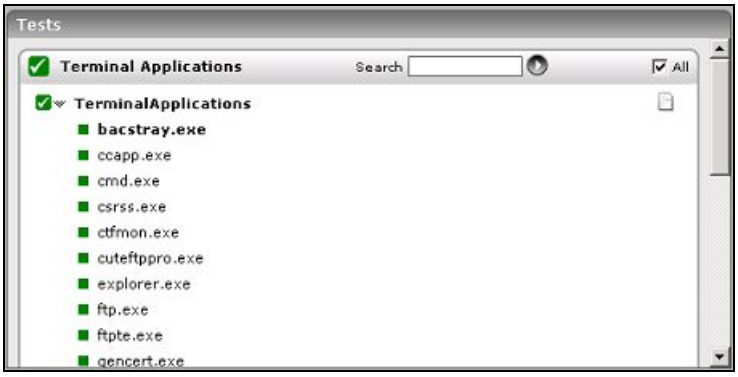

Figure 3.6: Tests associated with the Terminal Applications layer

3.4.1 Terminal Applications Test

This test reports statistics pertaining to the different applications deployed within the Microsoft RDS server and their usage by its clients.

Note:

This test will report metrics only if the Microsoft RDS server being monitored uses the .Net framework v3.5 (or above).

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results is reported for each application

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server.
Apps	By default, **all** is displayed here, which will auto-discover and monitor all the

Parameters	Description
	applications that are running from the Microsoft RDS server client. To monitor specific applications instead, you have to enter a comma separated list of processName:processPattern pairs which identify the applications published on the server being considered. 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, if a Microsoft Word application has been published on the Microsoft RDS server, then the **Process** to be specified is: Word:*winword*, where Word is the string to be displayed in the monitor interface, and *winword* is the application's executable. Other special characters such as slashes (\) can also be used while defining the process pattern. For example, if a server's root directory is /home/egurkha/apache and the server executable named httpd exists in the bin directory, then, the process pattern is "*/home/egurkha/apache/bin/httpd*". The test will rediscover the applications every 6th time the test runs.
Report by Domain Name	By default, this flag is set to **Yes**. This implies that by default, the detailed diagnosis of this test will display the *domainname\username* of each user who logged into the Microsoft RDS server. This way, administrators will be able to quickly determine which user logged in from which domain. If you want the detailed diagnosis to display the *username* alone, then set this flag to **No**.
Enable Browser Monitoring	By default, this flag is set to **No**, indicating that the eG agent does not monitor browser activity on the Microsoft RDS server. If this flag is set to **Yes,** then, whenever one/more IE (Internet Explorer) browser instances on the RDS server are accessed, the detailed diagnosis of the *Processes running* measure will additionally reveal the URL being accessed via each IE instance and the resources consumed by every URL. Armed with this information, administrators can identify the web sites that are responsible for excessive resource usage by an IE instance.
Show Active Apps Only	Using this flag, you can indicate whether the test should monitor all applications or applications that are currently active on the server. By default, this flag is set to **Yes**, indicating that only the currently active applications will be monitored by the eG agent. To monitor all applications, you need to set this flag to **No**.
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.

Parameters	Description
	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
Processes running	Number of instances of the published application currently executing on the Microsoft RDS server	Number	This value indicates if too many or too few instances corresponding to an application are executing on the host. The detailed diagnosis of this measure, if enabled, displays the complete list of processes executing, the users executing them, and their individual resource utilization.
CPU usage	Percentage of CPU used by the published application	Percent	A very high value could indicate that the specified application is consuming excessive CPU resources.
Memory usage	This value represents the ratio of the resident set size of the memory utilized by the application to the physical memory of the host system, expressed as a percentage.	Percent	A sudden increase in memory utilization for an application may be indicative of memory leaks in the application.
Handle count	Indicates the number of handles opened by this application.	Number	A consistent increase in the handle count over a period of time is indicative of malfunctioning of programs. Compare this value across applications to see which application is using a lot of handles.
Number of threads	Indicates the number of threads that are used by this application.	Number	

Measurement	Description	Measurement Unit	Interpretation
I/O data rate	Indicates the rate at which this application is reading and writing bytes in I/O operations.	KB/sec	This value counts all I/O activity generated by each instance of the application and includes file, network and device I/Os.
I/O data operations	Indicates the rate at which this application is issuing read and write data to file, network and device I/O operations.	Operations/sec	
I/O read data rate	Indicates the rate at which this application is reading data from file, network and device I/O operations.	KB/sec	
Page fault rate	Indicates the number of page faults generated by this application per second.	Faults/sec	Page faults are generated when an application tries to use memory that is part of its working set, but can't find it. Page faults can be either hard or soft: • Hard page faults occur when the page is found in the page file on the hard disk. The hard page faults involve disk I/O and impact performance • Soft page faults happen when the page is found somewhere else in memory. The soft page faults also impact performance, but may not result in heavy performance loss in a physical environment Compare values across the applications to figure out which application is generating most page faults.
Private working set	Indicates the total amount	MB	

Measurement	Description	Measurement Unit	Interpretation
memory	memory that is dedicated to this application, and cannot be shared among other applications on the server.		
Input delay for processes - max	Indicates the maximum amount of time lag detected between the user's input through any input device (e.g., mouse, keyboard) and the time at which this application detected the input.	Seconds	These measures will be reported only for the Windows 2019 desktops on the target server. Ideally, the values of these measures should be very low. High values for these measures can impact the speed of accessing the applications in the environment and badly degrade the overall user experience.
Input delay for processes - avg	Indicates the average amount of time lag detected between the user's input through any input device (e.g., mouse, keyboard) and the time at which this application detected the input.	Seconds	

The detailed diagnosis of the *Processes running* measure, if enabled, provides the list of processes currently executing, the users executing them, and their CPU and memory usage. Using these details, you can quickly detect resource-intensive instances and the user executing them.

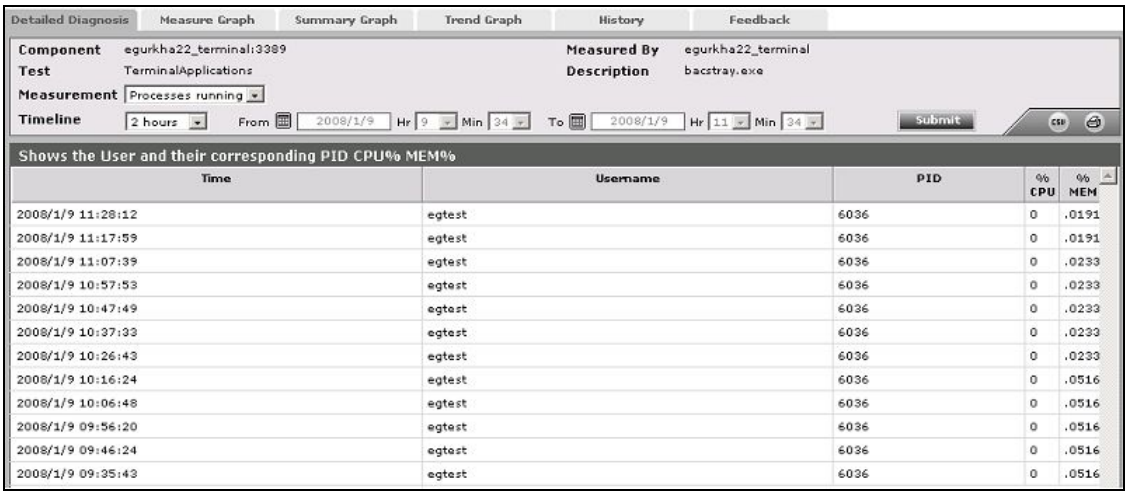

Figure 3.7: The detailed diagnosis of the Processes running measure

Moreover, if one or more browser instances are found to consume excessive CPU, memory and disk I/O resources on a server or a desktop, then for each such browser instance, administrators can now see a mapping of browser process to URL being accessed, as well as the resources used by each browser process in the detailed diagnosis. Armed with this information, administrators can determine the steps required to avoid excessive resource usage by browser instances – e.g., whether specific web sites are responsible for this, whether users are accessing web sites (e.g., youtube, facebook, etc.) that they should not be accessing from a corporate network, etc.

Note:

- The eG agent will perform browser activity monitoring only if the **ENABLE BROWSER MONITORING** flag is set to **Yes**.
- The eG agent will monitor browser activity only of the browser being accessed is **Internet Explorer**.

3.4.2 App-V Applications Test

This test reports statistics pertaining to the different applications executing on an App-V client and their usage. In addition, this test also reports the statistics pertaining to the processes running on the APP-V client.

Note:

This test will report metrics only when the App-V Client is installed on the Microsoft RDS Server.

Target of the test : An App-V Client on the target Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each application of the target App-V Client that is to be monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. By default, this is *NULL*.
Report By Domain Name	By default, this flag is set to **No**. This means that, by default, the test will report metrics for each username only. You can set this flag to **Yes**, to ensure that the test reports metrics for each domainname\username.
Extended Statistics	By default, this test provides you with detailed measures on the resource utilization of each application. If you wish to obtain only the CPU and memory related measures, then set this flag to **No**.
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
Total size	Indicates the total size	MB	The detailed diagnosis of this

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td>of this virtual application package.</td><td></td><td>measure lists the Version of the application, Application ID, Version ID of the applicaition and the application path.</td></tr>
<tr><td>Is loading?</td><td>Indicates whether this application is currently loading or not on the App-V client.</td><td></td><td>This measure reports a value True if the application is currently being loaded and a value False if otherwise.
These measure values and their corresponding numeric values are listed in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>True</td><td>1</td></tr><tr><td>False</td><td>0</td></tr></table>
Note:
By default, this measure reports the values Yes or No to indicate whether this application is currently being loaded on the client or not. The graph of this measure however is represented using the numeric equivalents - 0 or 1.</td></tr>
<tr><td>Loaded percentage</td><td>Indicates the percentage of this application that is currently being loaded on the App-V client.</td><td>Percent</td><td></td></tr>
<tr><td>In use?</td><td>Indicates whether this application is currently in use or not.</td><td></td><td>This measure reports a value True if the application is currently in use and a value False if otherwise.
These measure values and their</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>corresponding numeric values are listed in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>True</td><td>1</td></tr><tr><td>False</td><td>0</td></tr></table>
Note:
By default, this measure reports the values Yes or No to indicate whether this application is currently in use. The graph of this measure however is represented using the numeric equivalents - 0 or 1.</td></tr>
<tr><td>Any user based pending tasks available?</td><td>Indicates whether any tasks are pending for the user using this application.</td><td></td><td>This measure reports a value Yes if any tasks are pending for the user using the application and a value No if otherwise.
These measure values and their corresponding numeric values are listed 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 values Yes or No to indicate whether any tasks are currently pending for the user using this application. The graph of this measure however is represented</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>using the numeric equivalents - 0 or 1.</td></tr>
<tr><td>Any global based pending tasks available</td><td>Indicates whether any global tasks are pending for this application.</td><td></td><td>This measure reports a value Yes if any tasks are pending for the user using the application and a value No if otherwise.

These measure values and their corresponding numeric values are listed 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 values Yes or No to indicate whether any tasks are currently pending for the user using this application. The graph of this measure however is represented using the numeric equivalents - 0 or 1.</td></tr>
<tr><td>Processes running</td><td>Indicates the number of instances of this application currently executing.</td><td>Number</td><td>This value indicates if too many or too few instances corresponding to an application are executing on the host. The detailed diagnosis of this measure, if enabled, displays the complete list of processes executing, the users executing them, and their individual resource utilization.</td></tr>
<tr><td>CPU utilization</td><td>Indicates the</td><td>Percent</td><td>A very high value could indicate</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
	percentage of CPU used by this application.		that the specified application is consuming excessive CPU resources.
Memory utilization	This value represents the ratio of the resident set size of the memory utilized by the application to the physical memory of the host system, expressed as a percentage.	Percent	A sudden increase in memory utilization for an application may be indicative of memory leaks in the application.
Handle count	Indicates the number of handles opened by this application.	Number	An increasing trend in this measure is indicative of a memory leak in the process.
I/O data rate	Indicates the rate at which processes are reading and writing bytes in I/O operations.	Kbytes/Sec	This value counts all I/O activity generated by each process and includes file, network and device I/Os.
I/O data operations	Indicates the rate at which this application process is issuing read and write data to file, network and device I/O operations.	Operations/Sec	
I/O read data rate	Indicates the rate at which the process is reading data from file, network and device I/O operations.	Kbytes/Sec	
I/O write data rate	Indicates the rate at which the process is writing data to file, network and device I/O operations.	Kbytes/Sec	

Measurement	Description	Measurement Unit	Interpretation
Number of threads	Indicates the number of threads that are used by this application.	Number	
Page fault rate	Indicates the total rate at which page faults are occurring for the threads of all matching application processes.	Faults/Sec	A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. This may not cause the page to be fetched from disk if it is on the standby list and hence already in main memory, or if it is in use by another process with whom the page is shared.
Virtual memory used	Indicates the amount of virtual memory that is being used by the application.	MB	
Memory working set	Indicates the current size of the working set of a process.	MB	The Working Set is the set of memory pages touched recently by the threads in the process. If free memory in the computer is above a threshold, pages are left in the Working Set of a process even if they are not in use. When free memory falls below a threshold, pages are trimmed from Working Sets. If they are needed they will then be soft-faulted back into the Working Set before leaving main memory. If a process pattern matches multiple processes, the memory working set reported is the sum of the working sets for the processes that match the specified pattern. Detailed diagnosis for this

Measurement	Description	Measurement Unit	Interpretation
			test provides details of the individual processes and their individual working sets. Comparing the working set across processes indicates which process (es) are taking up excessive memory. By tracking the working set of a process over time, you can determine if the application has a memory leak or not.

3.4.3 Terminal Application Process Launches Test

To know which published applications on the Microsoft RDS server are currently accessed by users and how many instances of each application have been launched presently, use the **Terminal Application Process Launches** test. Detailed diagnostics, if enabled, reveal the users accessing the published applications and the thin clients from which the users are connecting to the Microsoft RDS server.

This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick the *Microsoft RDS* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the **<** button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every 'published application' on the Microsoft RDS server that is currently launched

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server .

Parameters	Description
Report by Domain Name	By default, this flag is set to **Yes**. This implies that by default, the detailed diagnosis of this test will display the *domainname\username* of each user who logged into the Microsoft RDS server. This way, administrators will be able to quickly determine which user logged in from which domain. If you want the detailed diagnosis to display the *username* alone, then set this flag to **No**.
Exclude	By default, this parameter is set to *none*. This means that the test will monitor all the applications that are launched on the Microsoft RDS server, by default. If you want the test to disregard certain applications when monitoring, then provide a comma-separated list of process names that correspond to the applications you want to ignore, in the exclude text box. For instance, your specification can be: **winword.exe,js.exe,taskmgr.exe**. Your specification can include wild card patterns as well. For example: ***win*,js*,*task**
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 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
Launch count	Represents the number of instances of this published application that have been launched currently.	Number	Use the detailed diagnosis of this measure to know which users are currently accessing the application and the clients from which the users are connecting.
Avg time to launch application	Indicates the average time taken by this application to launch.	Secs	Compare the value of this measure across applications to know which application took the longest time to launch. User experience with this application will naturally be poor.
Max time to launch application	Indicates the maximum time taken by this application to launch.	Secs	Compare the value of this measure across applications to know which application registered the highest launch time during the last measurement period. To know which user experienced this delay in launching, use the detailed diagnosis of the Launch count measure.

3.4.4 Outlook Add-ins Test

Outlook add-ins are integrations built by third parties into Microsoft Outlook using the new web technologies based platform. Microsoft Outlook add-ins have three key aspects:

- The same add-in and business logic works across desktop Microsoft Outlook for Windows and Mac, web (Office 365 and Outlook.com), and mobile.
- Outlook add-ins consist of a manifest, which describes how the add-in integrates into Outlook (for example, a button or a task pane), and JavaScript/HTML code, which makes up the UI and business logic of the add-in.
- Outlook add-ins can be acquired from the Office store or side-loaded by end-users or administrators.

The Outlook add-ins may be useful in connecting the business and social networks of the users. These add-ins when integrated with Microsoft Outlook simplifies the job of the users as they can stay up to date on the status and activities of their contacts by merely overlooking the Microsoft Outlook! When a user complains that it is taking too long to launch the add-ins of the Microsoft Outlook published on Microsoft RDS server, administrators must be able to quickly identify the add-ins that were loaded while the Microsoft Outlook is opened by the user, know how much time each add-in took to load, and thus pinpoint the add-in that is the slowest in loading. The **Outlook Add-ins** test provides these valuable insights to the administrators. This test auto-discovers all the add-ins integrated with the Microsoft Outlook published on the Microsoft RDS server, and for each discovered add-in, reports the number of times the add-in was loaded and the average and maximum time that add-in took to load. This way, the test points administrators to add-ins that are slow in loading.

This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the **<** button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every outlook add-in integrated with the Microsoft Outlook published on the Microsoft RDS server being monitored

Configurable parameters for the test

Parameters	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.
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 Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be

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

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Number of times loaded in last measure period	Indicates the number of times this outlook add-in was loaded during the last measurement period.	Number	The detailed diagnosis of this measure lists the time and duration for which the outlook add-in was loaded. Compare the value of this measure across the add-ins to figure out the most/least popular add-in.
Average load time	Indicates the average time taken by this outlook add-in to load.	Secs	
Maximum load time	Indicates the maximum time taken by this outlook add-in to load.	Secs	Compare the value of this measure across the add-ins to figure out the add-in that is the slowest to load.

3.5 The Terminal Users Layer

By continuously monitoring the user behavior on a Microsoft RDS server, administrators can accurately gauge resource usage per user, and derive guidelines for upgrading server capacity and imposing stricter access rules. The tests associated with this layer (see Figure 3.8) facilitate such user-related analysis.

Figure 3.8: Tests associated with the Terminal Users layer

3.5.1 Terminal Sessions Test

This test reports performance statistics related to Microsoft RDS server user sessions.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every server 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 the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server.
Ignore Down Session IDs	By default, this parameter is set to *65536,65537,65538* – these are nothing but the default ports at which the listener component listens. If any of these ports go down, then by default, this test will not count any of the sessions that failed when attempting to connect to that port as a Down session. You can override this default setting by adding more ports or by removing one/more existing ports.
Reportusingmanagertime	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**

Parameters	Description
	Diagnosis page for this test and in the **Thin Client** reports to be based on the Microsoft RDS server's local time.
Report by Domain Name	By default, this flag is set to **Yes**. This implies that by default, the detailed diagnosis of this test will display the *domainname**username* of each user who logged into the Microsoft RDS server. This way, administrators will be able to quickly determine which user logged in from which domain. If you want the detailed diagnosis to display the *username* alone, then set this flag to **No**.
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
Established sessions	Indicates the number of active terminal services sessions currently on the server.	Number	This measure gives an idea of the server workload in terms of active sessions. Tracking the number of active sessions with time, a Microsoft RDS server administrator can obtain information that can help him/her plan the capacity of their Microsoft RDS server farms. The detailed diagnosis capability, if enabled, lists the active and inactive sessions on

Measurement	Description	Measurement Unit	Interpretation
			the Microsoft RDS server.
Idle sessions	Indicates the number of sessions that are initialized and are currently ready to accept connections.	Number	To optimize the performance of a server, two default (idle) sessions are initialized before any client connections are made. For performance reasons, the number of idle sessions should be less than ten. Note that this test does not differentiate between RDP and ICA sessions.
Connected sessions	Indicates the current number of sessions that are connected, but no user has logged on to the server.	Number	A consistent increase in the value of this measure could indicate that users are having trouble logging in. Further investigation may hence be required. Note that this test does not differentiate between RDP and ICA sessions.
Connecting sessions	Indicates the number of sessions that are in the process of connecting.	Number	A very high value for this measure indicates a problem with the session or connection. Note that this test does not differentiate between RDP and ICA sessions.
Disconnected sessions	Indicates the number of sessions from which users have disconnected, but which are still active and can be reconnected.	Number	Too many disconnected sessions running indefinitely on a Microsoft RDS server cause excessive consumption of the server resources. To avoid this, a session limit is typically configured for disconnected sessions on the Microsoft RDS server. When a session limit is reached for a disconnected session, the session ends, which permanently deletes it from the server. Note that this test

Measurement	Description	Measurement Unit	Interpretation
			does not differentiate between RDP and ICA sessions.
Listen sessions	Indicates the current number of sessions that are ready to accept connections.	Number	Note that this test does not differentiate between RDP and ICA sessions.
Shadow sessions	Indicates the current number of sessions that are remotely controlling other sessions.	Number	A non-zero value for this measure indicates the existence of shadow sessions that are allowed to view and control the user activity on another session. Such sessions help in troubleshooting/resolving problems with other sessions under their control.
Down sessions	Indicates the current number of sessions that could not be initialized or terminated.	Number	Ideally, the value of this measure should be 0. By default, if sessions to any of these ports – 65536, 65537, 65538 – could not be initialized or terminated, they will not be counted as a 'down session'.
Init sessions	Indicates the current number of sessions that are initializing.	Number	A high value for this measure could indicate that many sessions are currently experiencing initialization problems.

The detailed diagnosis capability of the *Active sessions* measure, if enabled, lists the active and inactive sessions on the Microsoft RDS server, and provides details such as the user who initiated the sessions, the session login time, the duration for which the session was idle, etc.

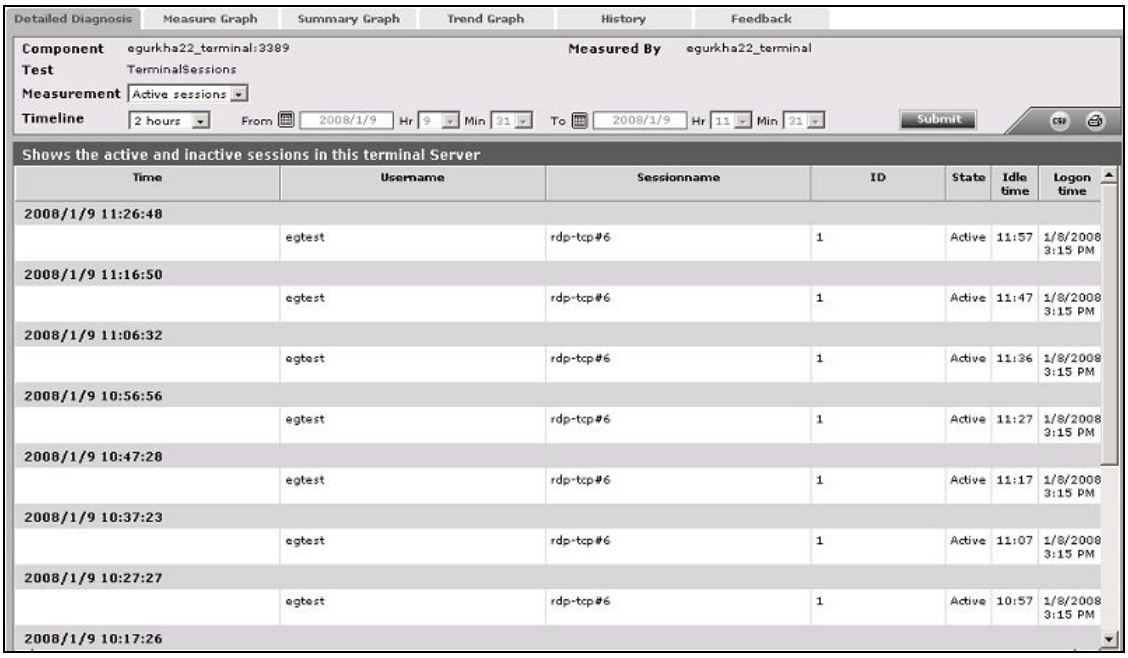

Figure 3.9: The detailed diagnosis of the Active sessions measure

3.5.2 Terminal Logins Test

This test monitors the new logins to the Microsoft RDS server.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results is reported for each Microsoft RDS server 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 the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server .
Reportusingmanagertime	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 Microsoft RDS server's local time.

Parameters	Description
Report by Domain Name	By default, this flag is set to **Yes**. This implies that by default, the detailed diagnosis of this test will display the *domainname\username* of each user session that logged out. This default setting ensures that administrators are able to quickly determine the domains to which the users who logged out belonged. You can set this flag to **No** if you want detailed diagnosis to display only the *username* of the users who logged out.
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 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
New logins	Indicates the number of new logins to the Microsoft RDS server in the last measurement period.	Number	A consistent zero value could indicate a connection issue.

Measurement	Description	Measurement Unit	Interpretation
Percent new logins	Indicates the percentage of current sessions that logged in during the last measurement period.	Percent	
Sessions logging out	Indicates the number of sessions that logged out.	Number	If all the current sessions suddenly log out, it indicates a problem condition that requires investigation.

The detailed diagnosis of the *Sessions logging out* measure lists the sessions that logged out.

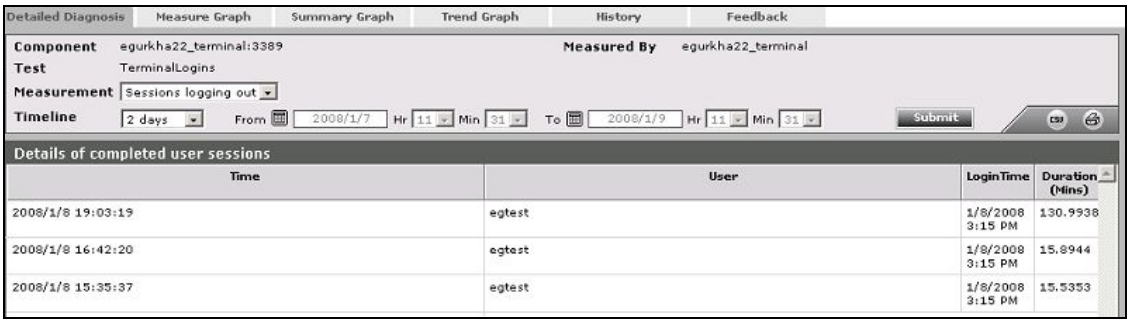

Figure 3.10: The detailed diagnosis of the Sessions logging out measure

3.5.3 Terminal Clients Test

This test measures the client connections to and from a Microsoft RDS server. This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick the *Microsoft RDS* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the **<** button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every server being monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.

Parameters	Description
Host	The host for which the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server .
ServerIP	By default, the ServerIP field will display the IP address of the Microsoft RDS server.
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
Current connections	The number of TCP connections currently established by clients to the Microsoft RDS server	Number	This measure directly indicates the loading on the Microsoft RDS server from clients. Typically one connection is established per active session to the Microsoft RDS server.
New connections	The number of new TCP connections initiated by clients to the Microsoft RDS server during the last measurement period	Number	Tracking the new connections over time can provide an indication of when clients login to the Microsoft RDS server. A spurt of connections and disconnections may be indicative of sporadic failures of the Microsoft RDS

Measurement	Description	Measurement Unit	Interpretation
			server.
Old connections removed	The number of TCP connections that were removed because the clients may have disconnected from the Microsoft RDS server during the last measurement period	Number	A large number of sudden connection drops may be early warning indicators of problems with the Microsoft RDS server.
Avg connection duration	The average time from when a connection is established to when the corresponding connection is disconnected. The duration of a connection is measured from its start time to the current time. The accuracy of this measurement is limited by the frequency at which this test is run.	Secs	This value can provide an indicator of how long clients stay connected to a Microsoft RDS server. This information together with the number of simultaneous clients can be useful for capacity planning in Microsoft RDS server environments (i.e., how to size the Microsoft RDS server).

3.5.4 Terminal Users Test

A Microsoft RDS server environment is a shared environment in which multiple users connect to a server/server farm and access a wide variety of applications. When server resources are shared, excessive resource utilization by a single user could impact the performance for other users. Therefore, continuous monitoring of the activities of each and every user on the server is critical. Towards this end, the **Terminal Users** test assesses the traffic between the user terminal and the server, and also monitors the resources taken up by a user's session on the server. The results of this test can be used in troubleshooting and proactive monitoring. For example, when a user reports a performance problem, an administrator can quickly check the bandwidth usage of the user's session, the CPU/memory/disk usage of this user's session as well as the resource usage of other user sessions. The admin also has access to details on what processes/applications the user is

accessing and their individual resource usage. This information can be used to spot any offending processes/ applications.

Note:

This test will report metrics only if the Microsoft RDS server being monitored uses the .Net framework v3.5 (or above).

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every user logged into the Microsoft RDS server

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server.
Usernames	Specify the name of the user whose performance statistics need to be generated. Multiple user names can be specified as a comma-separated list. all is used to indicate that all users of the Microsoft RDS server are to be monitored.
Report by Domain Name	By default, this flag is set to **Yes**. This implies that by default, this test will report metrics for every *domainname**username*. This way, administrators will know which user logged in from which domain. If you want the test to report metrics for every *username* only, then set this flag to **No**.
Enable Browser Monitoring	By default, this flag is set to **No**, indicating that the eG agent does not monitor browser activity on the Microsoft RDS server. If this flag is set to **Yes**, then, whenever one/more IE (Internet Explorer) browser instances on the RDS server are accessed, the detailed diagnosis of the **User sessions** measure will additionally reveal the URL being accessed via each IE instance and the resources consumed by every URL. Armed with this information, administrators can identify the web sites that are responsible for excessive resource usage by an IE instance.
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.

Parameters	Description
	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	Represents the current number of sessions for a particular user	Number	A value of 0 indicates that the user is not currently connected to the Microsoft RDS server.
CPU usage of user's processes	The CPU utilization for a session is the percentage of time that all of the threads/processes of a user session used the processor to execute instructions. If a user is connected via multiple sessions, the value reported is the sum of all cpu utilizations across all the sessions.	Percent	This value indicates the percentage of Cpu resources that are used by applications run by this user. Excessive CPU usage by a user can impact performance for other users. Check the detailed diagnosis to view the offending processes/applications.
Memory usage of user's processes	This value represents the ratio of the resident set size of the memory utilized by the user to the physical memory of the host system, expressed as a percentage. If a user is connected via multiple sessions, the value	Percent	This value indicates the percentage of memory resources that are used up by a specific user. By comparing this value across users, an administrator can identify the most heavy users of the Microsoft RDS server. Check the detailed diagnosis to view the offending processes/applications.

Measurement	Description	Measurement Unit	Interpretation
	reported is the sum of all memory utilizations across all the sessions.		
Input bandwidth	Indicates the average bandwidth used for client to server communications for all the sessions of a user	KB/Sec	This measure will not be available for Microsoft RDS servers running on Windows 2008 Service Pack 1 (or above).
Input errors	The average number of input errors of all types for all the sessions of a user. Example: Lost ACK's, badly formed packets, etc.	Errors/Sec	This measure will not be available for Microsoft RDS servers running on Windows 2008 Service Pack 1 (or above).
Output bandwidth	Indicates the average bandwidth used for server to client communications for all the sessions of a user	KB/Sec	This measure will not be available for Microsoft RDS servers running on Windows 2008 Service Pack 1 (or above).
Output errors	The average number of output errors of all types for all the sessions of a user. Example: Lost ACK's, badly formed packets, etc.	Errors/Sec	This measure will not be available for Microsoft RDS servers running on Windows 2008 Service Pack 1 (or above).
I/O read rate for user's processes	Indicates the rate of I/O reads done by all processes being run by a user.	KBps	These metrics measure the collective I/O activity (which includes file, network and device I/O's) generated by all the processes being executed by a user. When viewed along with the system I/O metrics reported by the DiskActivityTest, these measures help you determine the network I/O.

Measurement	Description	Measurement Unit	Interpretation
			Comparison across users helps identify the user who is running the most I/O- intensive processes. Check the detailed diagnosis for the offending processes/applications.
I/O write rate for user's processes	Indicates the rate of I/O writes done by all processes being run by a user.	KBps	
Faults for user's processes	Indicates the rate of page faults seen by all processes being run by a user.	Faults/Sec	Page Faults occur in the threads executing in a process. A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. If the page is on the standby list and hence already in main memory, or if the page is in use by another process with whom the page is shared, then the page fault will not cause the page to be fetched from disk. Excessive page faults could result in decreased performance. Compare values across users to figure out which user is causing most page faults.
Virtual memory of user's processes	Indicates the total virtual memory being used by all processes being run by a user.	MB	Comparison across users reveals the user who is being a drain on the virtual memory space.
Handles used by user's processes	Indicates the total number of handles being currently held by all processes of a user.	Number	A consistent increase in the handle count over a period of time is indicative of malfunctioning of programs. Compare this value across users to see which user is using a lot of handles. Check detailed diagnosis for further information.

Measurement	Description	Measurement Unit	Interpretation
CPU time used by user's sessions	Indicates the percentage of time, across all processors, this user hogged the CPU.	Percent	The CPU usage for user's processes measure averages out the total CPU usage of a user on the basis of the number of processors. For instance, if your Microsoft RDS server is using an 8-core processor and the total CPU usage of a user across all his/her sessions amounts to 80%, then the value of the CPU usage for user's processes measure for that user will be 10 % (80/8 processors = 10). This accurately denotes the extent of CPU usage in an environment where load is uniformly balanced across multiple processors. However, in environments where load is not well-balanced, the CPU usage for user's processes measure may not be an accurate indicator of CPU usage per user. For instance, if a single processor is used nearly 80% of the time by a user, and other 7 processors in the 8-core processor environment are idle, the CPU usage for user's processes measure will still report CPU usage as 10%. This may cause administrators to miss out on the fact that the user is actually hogging a particular processor! In such environments therefore, its best to use the CPU time used by user's sessions measure! By reporting the total CPU usage of a

Measurement	Description	Measurement Unit	Interpretation
			user across all his/her sessions and across all the processors the target Microsoft RDS server supports, this measure serves as the true indicator of the level of CPU usage by a user in dynamic environments. For instance, in the example above, the CPU time used by user's sessions of the user will be 80% (and not 10%, as in the case of the CPU usage for user's processes measure). A high value or a consistent increase in the value of this measure is hence serious and demands immediate attention. In such situations, use the detailed diagnosis of the CPU usage for user's processes measure to know what CPU-intensive activities are being performed by the user.
Input delay for user's sessions - max	Indicates the maximum amount of time lag detected between this user's input through any input device (e.g., mouse, keyboard) and the time at which the application detected the input.	Seconds	These measures will be reported only for the Windows 2019 desktops on the target RDS server. Ideally, the values of these measures should be very low. High values for these measures can impact the speed of accessing the applications in the environment and seriously degrade the overall user experience.
Input delay for user's sessions - avg	Indicates the average amount of time lag detected between this user's input through any input device (e.g.,	Seconds	

Measurement	Description	Measurement Unit	Interpretation
	mouse, keyboard) and the time at which the application detected the input.		

The detailed diagnosis of the *User sessions, CPU usage of user's processes*, and *Memory usage of user's processes* measures lists the processes executed by a user on the Microsoft RDS server, and reports the resource usage of each process (see Figure 3.11).

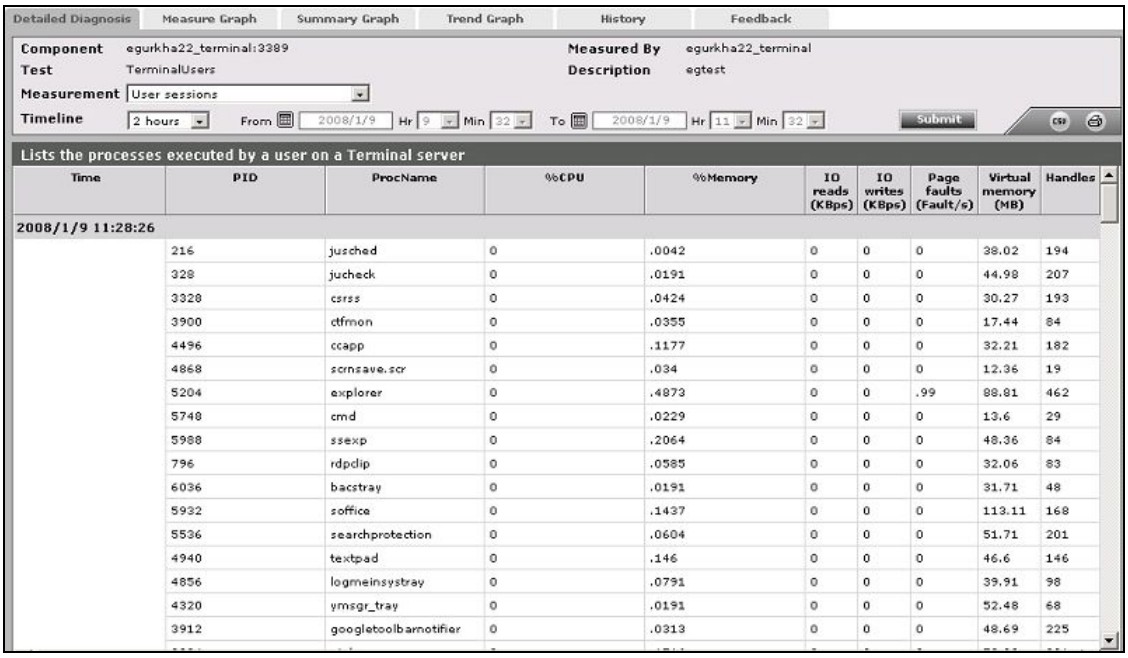

Time	PID	ProcName	%CPU	%Memory	IO reads (KBps)	IO writes (KBps)	Page faults (Fault/s)	Virtual memory (MB)	Handles
2008/1/9 11:28:26									
	216	jusched	0	.0042	0	0	0	38.02	194
	328	jucheck	0	.0191	0	0	0	44.98	207
	3328	csrss	0	.0424	0	0	0	30.27	193
	3900	ctfmon	0	.0355	0	0	0	17.44	84
	4496	ccapp	0	.1177	0	0	0	32.21	182
	4868	scrnsave.scr	0	.034	0	0	0	12.36	19
	5204	explorer	0	.4873	0	0	.99	88.81	462
	5748	cmd	0	.0229	0	0	0	13.6	29
	5988	ssexp	0	.2064	0	0	0	48.36	84
	796	rdpclip	0	.0585	0	0	0	32.06	83
	6036	bacstray	0	.0191	0	0	0	31.71	48
	5932	soffice	0	.1437	0	0	0	113.11	168
	5536	searchprotection	0	.0604	0	0	0	51.71	201
	4940	textpad	0	.146	0	0	0	46.6	146
	4856	logmeinsystray	0	.0791	0	0	0	39.91	98
	4320	ymsgr_tray	0	.0191	0	0	0	52.48	68
	3912	googletoolbarnotifier	0	.0313	0	0	0	48.69	225

Figure 3.11: The detailed diagnosis of the User sessions measure

Where one/more instances of the Internet Explorer browser are running, the detailed diagnosis additionally displays the website URL accessed using each IE instance, the domain of every URL, and the website title. In the event of excessive resource usage by an IE instance, this information will shed light on the resource-intensive web site that was being accessed.

Note:

- The eG agent will perform browser activity monitoring only if the Enable Browser Monitoring flag is set to **Yes**.
- The eG agent will monitor browser activity only of the browser being accessed is **Internet Explorer**.

3.5.5 Terminal Users By Browsers Test

In recent times, browsers have become one of the common ways to access many applications in enterprises. The same browser may be used for accessing multiple applications. Further, users may even use browsers to access non- corporate web sites from their environment. In modern architectures, a lot of the processing is done by the scripts executed on browsers. This further adds processing tasks to the browsers. Virtual server administrators need to know when exactly specific browser instances started taking up excessive resources (CPU, memory, disk) and most importantly, what URLs were accessed when a browser started taking up resources. The **Terminal Users By Browsers** test helps administrators in this regard!

This test auto-discovers the browsers accessed by each user connected to a Microsoft RDS farm, and for each browser, reports the number of sessions initiated and the number of processes running. This test also reports the resource utilization of each browser. Using this test, administrators can figure out the browser that is responsible for excessive resource utilization of a user.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal/remote agent

Outputs of the test : One set of results for the Microsoft RDS server being monitored

Configurable parameters for the test

Parameters	Description
Test Period	How often should the test be executed. By default, this is 15 minutes.
Host	The host for which the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server.
Report by Domain Name	By default, the flag is set to **Yes**. This implies that by default, this test will report metrics for every domainname\username configured for this test. This way, administrators will be able to quickly determine which user logged in from which domain. If you want the test to report metrics for the username alone, then set this flag to **No**.
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

Parameters	Description
	Frequency.
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
User sessions	Indicates the total number of user sessions initiated through this browser.	Number	
Processes running in user's sessions	Indicates the number of processes running for this browser.	Number	The detailed diagnosis of this measure lists the Process ID, Website Title, Website URL and the Website domain.
CPU usage for user's processes	Indicates the percentage of CPU utilized by the processes running for this browser.	Percent	A high value for this measure is a cause of concern. Comparing the value of this measure across browsers helps administrators in identifying the browser that is utilizing too much of CPU resources. The detailed diagnosis of this measure lists the session ID, Process ID, the CPU utilized by the process and the memory utilized by the process.
Memory usage for user's processes	Indicates the percentage of memory utilized by the	Percent	A high value for this measure is a cause of concern.

Measurement	Description	Measurement Unit	Interpretation
	processes running for this browser.		Comparing the value of this measure across browsers helps administrators in identifying the browser that is utilizing too much of memory resources. The detailed diagnosis of this measure lists the session ID, Process ID, the CPU utilized by the process and the memory utilized by the process.

3.5.6 Terminal Disconnects Test

A user session is terminated when a user logs off from the Citrix/Microsoft RDS server or when the session is abruptly interrupted (e.g., due to server, network, or application errors). When a user logs off, all the applications started by the user are terminated. However, when a user disconnects, the applications started by the user will keep running on the server consuming resources. Hence, the number of disconnected sessions on a Citrix/Microsoft RDS server should be kept to a minimum. Abrupt disconnects can significantly impact the end user experience, and hence, it is important to monitor the number of disconnected sessions at any point of time.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results is reported for each Microsoft RDS server 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 the test is to be configured.
Port	Refers to the port used by the Microsoft RDS server .
Reconnectperiod	This parameter is used by the test while computing the value for the *Quick reconnects* measure. This measure counts all the users who reconnected to the Microsoft RDS server within the short period of time (in minutes) specified against Reconnectperiod.

Parameters	Description
Report by Domain Name	By default, this flag is set to **Yes**. This implies that by default, the detailed diagnosis of this test will display the *domainname\username* of each user who logged into the Microsoft RDS server. This way, administrators will be able to quickly determine which user logged in from which domain. If you want the detailed diagnosis to display the *username* alone, then set this flag to **No**.
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 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
Total disconnected sessions	Indicates the total number of sessions that are in the disconnected state.	Number	
New disconnects	Indicates the number of sessions that were disconnected in the last	Number	The detailed diagnosis of this measure, if enabled lists the users who have recently disconnected.

Measurement	Description	Measurement Unit	Interpretation
	measurement period		
Quick reconnects	Indicates the number of users who reconnected soon after a disconnect.	Number	The detailed diagnosis of this measure, if enabled lists the users who have reconnected quickly.

The detailed diagnosis for the *New disconnects* measurement indicates the user, session ID, and client type for each newly disconnected session. This information can be used to track whether specific users are being disconnected often (see Figure 3.12).

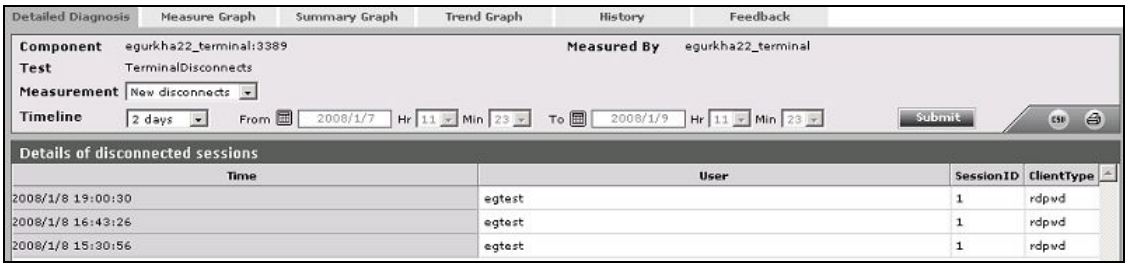

Figure 3.12: The detailed diagnosis of the New disconnects measure

The detailed diagnosis for the *Quick reconnects* measurement indicates the user, session ID, client type, the exact time at which the session disconnected, and duration of the disconnect, for each session that quickly reconnected. This information can be used to track whether specific users are being disconnected often (see Figure 3.13).

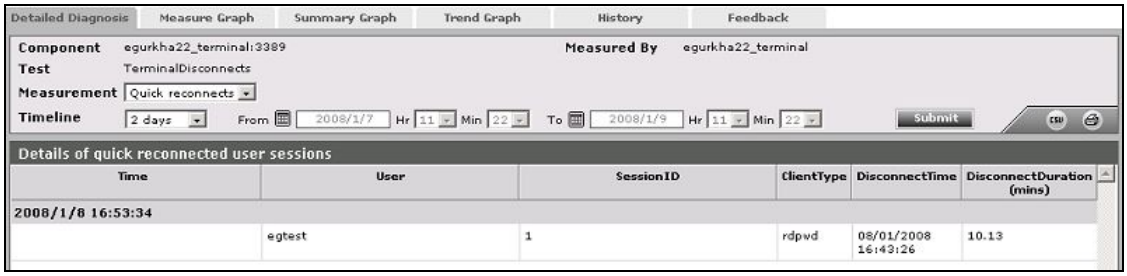

Figure 3.13: The detailed diagnosis of the Quick reconnects measure

3.5.7 Rdp Client Access Test

A Microsoft RDS server environment is a shared environment in which multiple users connect to a server from remote terminals using the Remote Desktop Protocol (RDP). One of the key factors influencing user experience in such an environment is the latency seen by the users when connecting to the server. High network latencies or packet losses during transmission can cause significant slow-downs in request processing by the server. Hence, monitoring latencies between the server and individual client terminals is important.

The Rdp Client Access test is executed by the eG agent on a Microsoft RDS server. This test auto-discovers the users who are currently logged on to the server and the IP address from which they are connecting to the Microsoft RDS server. For each user, the test monitors the quality of the link between the client and the Microsoft RDS server.

Using this test, an administrator can identify user sessions that are being impacted by high latencies or by excessive packet drops. In some cases, a Microsoft RDS server may regard a user session as active, even though the network link connecting the user terminal to the Microsoft RDS server has failed. The Rdp Client Access test alerts administrators to such situations.

This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick the Microsoft RDS ad the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the **<** button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of outputs for every user currently connected to the Microsoft RDS server

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the host listens
Displaydomain	By default, the *DISPLAYDOMAIN* flag is set to **Yes;** this indicates that thistest, by default, will report metrics for every **domainname\username** who is currently connected to the server. This way, administrators can quickly figure out which

Parameters	Description
	user is connecting to the server from which domain. You can set this flag to **No** to ensure that this test reports metrics for each **username** only.
Packetsize	The size of packets used for the test (in bytes)
Packetcount	The number of packets exchanged between the Microsoft RDS server and the user terminal during the test
Timeout	How long after transmission should a packet be deemed lost (in seconds)
PacketInterval	Represents the interval (in milliseconds) between successive packet transmissions during the execution of this test.
Reportunavailability	By default, this flag is set to **No**. This implies that, by default, the test will not report the unavailability of network connection between a user terminal and the Microsoft RDS server. In other words, if the **Packet loss** measure of this test registers the value **100%**for any user, then, by default, this test will not report any measure for that user; under such circumstances, the corresponding user name will not appear as a descriptor of this test. You can set this flag to **Yes**, if you want the test to report and alert you to the unavailability of the network connection between a user terminal and the Microsoft RDS server.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Number of sessions	Indicates the current number of sessions for a particular user	Number	The value 0 indicates that the user is not currently connected to the Microsoft RDS server.
Average delay	Indicates the average delay between transmission of a request by the agent on a Microsoft RDS server and receipt of the response back from the user terminal.	Secs	Comparing the value of this measure across users will enable administrators to quickly and accurately identify users who are experiencing higher latency when connecting to a Microsoft RDS server.
Minimum delay	Indicates the minimum	Secs	A significant increase in the

Measurement	Description	Measurement Unit	Interpretation
	delay between transmission of a request by the agent on a Microsoft RDS server and receipt of the response back from the user terminal.		minimum round-trip time is often a sure sign of a poor link between the server and a user's terminal.
Packet loss	Indicates the percentage of packets lost during data exchange between the Microsoft RDS server and the user terminal.	Percent	Comparing the value of this measure across users will enable administrators to quickly and accurately identify users who are experiencing slowdowns because of poor performance on the network links between their terminals and the Microsoft RDS server.

Note:

- If the same user is connecting to the Microsoft RDS server from multiple client terminals, the value of the *Number of sessions*, *Avg delay*, and *Packet loss* measures will be averaged across all the sessions of that user. The *Minimum delay* measure, on the other hand, will display the least value reported for *Minimum delay* across all the sessions of that user.
- When a user logs out, the number of sessions will be reduced by 1. If the number of user sessions becomes 0, the corresponding entry for that user in the eG user interface will be removed after a short period of time.
- By default, the Rdp Client Access test spawns a maximum of one thread at a time for monitoring each of the RDP connections to a Microsoft RDS server. Accordingly, the **MaxRdpClientThreads** parameter in the **eg_ tests.ini** file (in the <EG_ INSTALL_ DIR>\manager\config directory) is set to 1 by default. In large Microsoft RDS server farms however, numerous concurrent users attempt to connect to the Microsoft RDS server from multiple remote client terminals. To enhance the efficiency of the test and to make sure that it scales to monitor the large number of RDP connections to the Microsoft RDS server, you might want to consider increasing the value of the **MaxRdpClientThreads** parameter. If this parameter is set to say, 15, then, it implies that the test will spawn a maximum of 15 threads at one shot, thus monitoring 15 RDP connections to the Microsoft RDS server, simultaneously.

3.5.8 RemoteFX User Experience Test

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

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

Note:

This test works only on Windows 2008 Service Pack 1 (or above).

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every user logged into the Microsoft RDS server

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the host listens
Report by Domain Name	By default, this flag is set to **Yes**. This implies that by default, this test will report metrics for every *domainname**username*. This way, administrators will know which user logged in from which domain. If you want the test to report metrics for every *username* only, then set this flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
User sessions	Represents the current number of sessions for a particular user.	Number	A value of 0 indicates that the user is not currently connected to the Microsoft RDS server.
Average frames encoding time	Indicates the average time taken for encoding the frames of this user.	Secs	Compare the value of this measure across users to know for which user frames encoding took too long.
Frame quality	Indicates the quality of the output frame expressed as a percentage of the quality of the source frame for this user.	Percent	High frame rates produce a smooth representation of frames for the particular user, while low frame rates may cause rough or choppy representation of frames for the particular user. A high value is hence desired for this measure. Compare the value of this measure across users to know which user received the poorest frame quality.
Frames skipped due to insufficient client resources	Indicates the rate at which frames were skipped for this user due to insufficient client resources.	Frames/Sec	A low value is desired for this measure. Compare the value of this measure across users to know which user is connecting from a client sized with inadequate resources.
Frames skipped due to insufficient network resources	Indicates the rate at which frames were skipped for this user due to insufficient network resources.	Frames/Sec	A low value is desired for this measure. Compare the value of this measure across users to know which user is connecting via a network that is sized with inadequate resources.
Frames skipped due to insufficient server resources	Indicates the rate at which frames were skipped for this user due to insufficient server resources.	Frames/Sec	A low value is desired for this measure. Compare the value of this measure across users to know which user was unable to receive frames due to the lack of enough resources on the Microsoft RDS server.
Graphics compression ratio	Indicates the ratio of the number of bytes encoded to the number of bytes	Percent	The compression ratio typically affects the quality of the picture. Generally, the higher the compression ratio, the

Measurement	Description	Measurement Unit	Interpretation
	input for this user.		poorer the quality of the resulting picture. Ideally therefore, the value of this measure should be 0. You can compare the value of this measure across users to identify that user whose picture output was very poor owing to high compression.
Input frames	Indicates the number of source frames provided per second as input to the RemoteFx graphics for this user.	Frames/Sec	
Output Frames	Indicates the number of source frames sent per second to this user as output of RemoteFx graphics.	Frames/Sec	
Source frames	Indicates number of frames per second composed at the source for this user.	Frames/Sec	
Base TCP round trip time	Indicates the time between initiating a network request and receiving a response over TCP for this user.	Secs	A high value for this measure could indicate a bottleneck in TCP connectivity between the user terminal and the server.
Base UDP round trip time	Indicates the time between initiating a network request and receiving a response over UDP for this user.	Secs	A high value for this measure could indicate a bottleneck in UDP connectivity between the user terminal and the server.
Current TCP bandwidth	Indicates the amount of data that is currently carried from one point to another over TCP for this user.	Kbps	A consistent rise in the value of this measure could indicate that TCP traffic to/from the user is consuming bandwidth excessively. Compare the value of this measure across users to identify that user who is performing

Measurement	Description	Measurement Unit	Interpretation
			bandwidth-intensive operations on the Microsoft RDS server.
Current TCP round trip time	Indicates the average time between initiating a network request and receiving a response over TCP for this user.	Secs	A high value could indicate a current problem with TCP connectivity between the user terminal and the server.
Current UDP bandwidth	Indicates the amount of data that is currently carried from one point to another over UDP for this user.	Kbps	A consistent rise in the value of this measure could indicate that UDP traffic to/from the user is consuming bandwidth excessively. Compare the value of this measure across users to identify that user who is performing bandwidth-intensive operations on the Microsoft RDS server.
Current UDP round trip time	Indicates the average time between initiating a network request and receiving a response over UDP for this user.	Secs	A high value could indicate a current problem with UDP connectivity between the user terminal and the server.
Forward error correction rate	Indicates the percentage of forward error corrections performed for this user.	Percent	RemoteFX UDP transport uses Forward Error Correction (FEC) to recover from the lost data packets. In the cases where such packets can be recovered, the transport doesn't need to wait for the data to be retransmitted, which allows immediate delivery of data and prevents Head of Line Blocking. Preventing this stall results in an overall improved responsiveness. A high value is hence desired for this measure.
Loss	Indicates the percentage of packets lost when being transmitted to this user.	Percent	A high value indicates that a large number of packets were lost without being retransmitted. By comparing the value of this measure across users, you can find that user who has suffered

Measurement	Description	Measurement Unit	Interpretation
			the maximum data loss. This could be owing to a bad network connection between the remote user terminal and the server.
Retransmission	Indicates the percentage of packets that have been retransmitted to this user.	Percent	Retransmissions should only occur when it is certain that a packet to be retransmitted was actually lost. Redundant retransmissions can also occur because of lost acknowledgments, coarse feedback, and bad retransmissions. Retransmission rates over 5% can indicate degraded network performance on a LAN. The internet may vary between 5 and 15 percent depending upon traffic conditions. Any value above 25 percent indicates an excessive number of retransmissions that will significantly increase the time for the file transfer and annoy the user.
TCP received rate	Indicates the rate at which the data is received over TCP for this user.	Kbps	A high value is desired for these measures as it indicates high TCP throughput.
TCP sent rate	Indicates the rate at which the data is sent over TCP for this user.	Kbps	
UDP received rate	Indicates the rate at which the data is received over UDP for this user.	Kbps	A high value is desired for these measures as it indicates high UDP throughput.
UDP sent rate	Indicates the rate at which the data is sent over UDP for this user.	Kbps	

Note:

Optionally, you can enable an **EventLog** test for the Microsoft RDS server to closely monitor the system and application events on the server. This test is disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests ->

Enable/Disable, pick the *Microsoft RDS* as the desired **Component type**, set *Performance* as the **Test type**, choose the test from the **DISABLED TESTS** list, and click on the **<** button to move the test to the **ENABLED TESTS** list. Finally, click the **Update** button. This test is mapped to the **Windows Service** layer of the Microsoft RDS server component.

3.5.9 Terminal Server Input Delay Test

One of the most difficult issues faced by users during the RDP sessions is slow response from applications. Every time, the slowness might not be caused due to network latency, workload, a hardware configuration issue or an operating system issue. In some cases, the users might experience slowness due to input device lag that is the time delay experienced between the user's input via a keyboard or mouse and the time at which the application responds to the input. This input delay can deny quick access to the applications and significantly degrade the overall user experience during RDP sessions. Whenever the users experience slowness in accessing the applications, administrators need to figure out if a delay is caused due to the network latency, workload, hardware configuration issue, operating system issue or input device lag. By knowing what is causing the delay, administrators can effectively troubleshoot the problem condition and resolve them in time, so that the user experience during the RDP sessions will not get impacted. To help administrators in this regard, eG offers the **Terminal Server Input Delay** test.

This test monitors the RDS server and reports the maximum and average time taken by the applications to respond to the user input across all RDP sessions. Using this revelation, administrators can easily figure out if the slowness is caused due to the input device lag and, if so, take remedial actions to resolve the issue before it seriously affects the user experience.

Note:

This test will report measures only for the Windows 2019 desktops on the Microsoft RDS server.

Target of the test : A Microsoft RDS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the target RDS server being monitored.

Configurable parameters for the test

1. **TEST PERIOD** – How often should the test be executed
2. **HOST** – The host for which the test is to be configured
3. **PORT** – Refers to the port used by the Microsoft RDS server.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Maximum input delay for sessions	Indicates the maximum amount of time lag detected between the user's input through any input device (e.g., mouse, keyboard) and the time at which the application responds to the input.	Seconds	Ideally, the values of these measures should be zero or very low. High values for these measures can impact the speed of accessing the applications in the environment which in turn seriously degrades the overall user experience. To know exactly which user/application experiences the maximum input lag, you can refer to the **Terminal Users** and **Terminal Applications** tests.
Average input delay for sessions	Indicates the average amount of time lag detected between the user's input through any input device (e.g., mouse, keyboard) and the time at which the application detected the input.	Seconds	

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.